



Authentication, Authorization, and Accounting

Alessandro Paolini¹ , Diego Scardaci¹, Nicolas Liampotis² ,
Vincenzo Spinoso¹ , Baptiste Grenier¹ , and Yin Chen¹ 

¹ EGI Foundation, Amsterdam, The Netherlands
{alessandro.paolini,diego.scardaci,vincenzo.spinoso,
baptiste.grenier,yin.chen}@egi.eu
² GRNET, Athens, Greece
nliam@grnet.gr

Abstract. Environmental research infrastructures and data providers are often required to authenticate researchers and manage their access rights to scientific data, sensor instruments or online computing resources. It is widely acknowledged that Authentication, Authorization and Accounting (AAA) play a crucial role in providing a secure distributed digital environment. This chapter reviews the advanced AAA technology and best practices in the existing pan-European e-Infrastructures. It also discusses the challenging issues of interoperability in federated access and presents state-of-the-art solutions.

Keywords: Authentication · Authorization · Accounting

1 Introduction

A challenge that any operational Research Infrastructure (RI) has to deal with is controlling the access to these services and resources: the identity of the users needs to be verified, and once this is done successfully, the proper rights have to be granted to the users to perform the operations they are supposed to do. It is widely acknowledged that Authentication, Authorization and Accounting (AAA) play a crucial role in providing a secure distributed digital environment. In the rest of the chapter, we will discuss the first two ‘AA’s – Authentication and Authorization; then, address the issues for the last ‘A’ – Accounting, separately. We will review the state-of-the-art of AAA, and discuss the best practice in EGI e-Infrastructure [1]¹.

All the procedures, policies, and technologies used to implement such basic activities are part of the so-called Authentication and Authorization Infrastructure (AAI): it is a key service meant to ensure that services and resources are accessed by the users in a secure way and that at the same time the users personal data are stored in a safe manner.

¹ EGI is the first European-wide publicly-funded e-Infrastructure. It currently federates 237 computing centres across Europe and world-wide, providing high-throughput (grid and cloud) computing, storage and data resources to support European research, at the moment, having over 1 Million CPU cores and almost 700 Petabytes (disk and tape) storages.

Different AAI technologies have been using over the years to secure access to digital devices, the most common practice in e-Infrastructures such as EGI is the Public Key Infrastructure (PKI), which will be outlined in Sect. 2.

While PKI technology has been generally used for the access mainly to non-web based services, there was also the necessity to assign a digital (and single) identity to the users in order to regulate the access (generally) to web-based services. The need for user identity to cross borders between organisations, domain and services, lead to the creation of federated identity environments. Home organisations (e.g. a university, library or research institute.), who operate an Identity Provider (IdP)², register users by assigning a digital identity – in this way, they are able to authenticate their users and provide a limited set of attributes that characterise the user in a given context. Resource owners (Service Providers) delegate the authentication to Identity Providers in order to control access to the provided resources. An Identity federation is a group of Identity and Service Providers that sign up to an agreed set of policies for exchanging information about users and resources to enable access to and use of the resources. There are many Research and Education identity federations around the globe and they commonly have a national coverage³: for example, eduGAIN [2] interconnects identity federations around the world, simplifying access to content, services and resources for the global research and education community. We are reporting about issues, challenges and requirements for operating interoperable AAIs in Sect. 3.

In Sect. 4 we will discuss the solution to the identify federations and depict the “AARC Blueprint Architecture (BPA)⁴”, created with the purpose to provide a set of interoperable architectural building blocks for software architects and technical decision-makers, who are designing and implementing access management solutions for international research collaborations.

In order to provide an implementation example of AARC BPA, in Sect. 5, we will describe the Check-in service, the AAI platform for the EGI Infrastructure.

Section 6 is about Accounting (the third “A” of AAA). Accounting provides the method for collecting and sending user activity information used for billing, auditing, and reporting, such as user identifies, start and stop times, executed commands, number of packets, and number of bytes. The accounting tool used by EGI is called APEL, originally created for the LHC Computing Grid (LCG). APEL parses batch, system and gatekeeper logs generated by a site and builds accounting records, which provide a summary of the resources consumed based on the attributes, such as CPU time, Wall Clock Time, Memory and EGI user DN (Domain Name). APEL is the underpinning technology of the EGI accounting portal [4] that supports the daily operation of e-Infrastructure.

Finally, we will conclude this chapter in Sect. 7.

² By definition, an IdP is a system that creates, maintains, and manages identity information for principals (users, services, or systems) and provides principal authentication to other service providers (applications) within a federation or distributed network.

³ The REFEDs map to discover worldwide identity federations <https://refeds.org/federations/federations-map>.

⁴ AARC Blueprint Architecture (BPA) [16] is produced by the AARC project [3], an EC H2020 project, aims to address the increased need for federated access and for authentication and authorisation mechanisms by research and e-infrastructures.

2 Public Key Infrastructure and Digital Certificates

The purpose of this section is giving an overview of the basilar concepts of the Public Key Infrastructure, providing the required information useful for the context of this document: for a more exhaustive description, the reader can have a look at the references linked to this section.

A Public Key Infrastructure (PKI) is a set of roles, policies, procedures and technologies to authenticate electronic users and devices: by using a cryptographic technique it enables entities to securely communicate on an insecure public network, and reliably verify the identity of an entity via digital signatures [5]. The identity of each entity is bound to a key pair to encrypt and decrypt messages: a public and a private key which are mathematically related. Therefore, this model makes use of asymmetric cryptography algorithms with the following properties:

1. It is impossible to derive the private key from the public one.
2. The public key can be distributed to other entities in the system to encrypt messages which can be decrypted only by the corresponding private key, which therefore must be kept secret.

Let's assume that two entities in the system, John and Beth, need to exchange some digital content between them. The following basilar steps will be accomplished:

1. Both John and Beth have their own key pairs. They are safely storing their private key and they have sent their public key to each other.
2. Before sending a message to Beth, John encrypts it using her public key.
3. To decrypt the message, Beth uses her private key.

Anyway, this simplified process highlights an important concern: how can John be sure that the public key used to encrypt the message for Beth really belongs to her?

To address this issue, one way to do is to introduce in the model a trusted third party that certifies the integrity and the ownership of the public keys: this new entity is called Certification Authority (CA) and has the important role of storing, issuing and signing the digital certificates used to verify that a particular public key belongs to a certain entity. The CA certificate could be self-signed or signed by another CA, and it is used to sign all the Certificate Signing Request (CSR) containing public keys. At the same time, a CA periodically publishes the so-called Certificate Revocation List (CRL) which contains a list of all the revoked certificates: this is to constantly make aware all the entities about the validity of the issued certificates.

When a CA is used, the previous example can be modified in the following way:

1. Assume that the CA has issued a digital certificate that contains its public key. This certificate is signed with the CA private key.
2. Beth and John agree to use the CA to verify their identities.
3. Beth requests a certificate to the CA, by sending it a CSR containing her public key.
4. The CA verifies her identity and issues the certificate making it publicly available.

5. John retrieves the certificate, verifies it, so he can assume that the public key in the certificate does indeed belong to Beth.
6. John uses Beth's verified public key to encrypt a message to her.
7. Beth uses her private key to decrypt the message from Bob.

Another use case is when John wants to send a message to Beth allowing Beth to verify that the message has really been sent by John. In such a case, John should digitally sign the message with his private key and Beth can verify his identity validating the signature with John's public key. This case is particularly relevant for distributed infrastructure, including EGI, because it enables EGI services to verify the identity (the authentication process) of a user submitting a task (e.g. run a workflow in the EGI infrastructure).

In summary, digital certificates are a way to perform mutual authentication between two parties: in this process, the two parties authenticate each other through verifying the provided digital certificate so that both parties are assured of the others' identity. In technology terms, it refers to a client (web browser or client application) authenticating themselves to a server (website or server application) and that server also authenticating itself to the client through verifying the public key certificate/digital certificate issued by the trusted CAs. The process is therefore called "certificate-based mutual authentication" [8]: since the users can securely access a server by exchanging a digital certificate instead of a username and password, this helps in preventing phishing, keystroke logging and man-in-the-middle (MITM) attacks among other common problems with password-based authentication.

The current standard that defines digital certificates is called X.509 Version 3 [7]. An X.509 v3 certificate includes the following elements:

- Certificate serial number
- The digital signature of the CA
- The public key of the user to whom the certificate is issued
- Identity of the owner
- Date of expiration
- Name of the CA that has issued the certificate

In the digital world, one single CA usually covers a predefined geographic region or administrative domain (if not all the world), such as an organization, a country, or a set of countries, therefore the identity vetting process would not scale-up if done by the CA itself. This is the reason why the task of verifying the identity (and personal data) of entities requesting their digital certificates is usually delegated to a network of subordinated Registration Authorities (RAs) who act on behalf of the parent CA in their assigned sub-domain.

In a world where large scale distributed computing is deployed on a production scale, across organisations, across countries, and across continents, a common trust domain for distributed computing has been created to join the several existing certification authorities into a single authentication domain and thus enabling sharing of computing and resources worldwide: the Interoperable Global Trust Federation (IGTF) [6] has been created to coordinate and manage this trust domain. The IGTF is a body to establish common policies and guidelines that help establish interoperable, global trust relations between

providers of e-Infrastructures and cyber-infrastructures, identity providers, and other qualified relying parties. It is divided in three Policy Management Authorities (PMAs) covering the Asia Pacific, the Americas and Europe, Middle-East and Africa.

2.1 Proxy Delegation

In order to use large scale distributed computing infrastructure, such as EGI, a user need a way to copy its own credentials to the machines where its workflows/jobs are going to be executed: this is necessary to allow remote sub-processes or other services in the e-infrastructure to perform on behalf of the user (delegation), particular operations needed to successfully complete the workflow, like the access of data belonging to the user stored on different resource providers, or start sub-jobs on other resources. If the user really utilises their own long-living personal certificate to do so, this would lead to security problems (in case of stolen credentials), and it would also make difficult arranging in advance the several delegations: that is why the X.509 Proxy Certificates [9] have been created.

Proxy Certificates allow an entity holding a standard X.509 public-key certificate to delegate some or all of its privileges to another entity which may not hold X.509 credentials at the time of delegation. This delegation can be performed dynamically, without the assistance of a third party, and can be limited to arbitrary subsets of the delegating entity's privileges. Once acquired, a Proxy Certificate is used by its bearer to authenticate and establish secure connections with other parties in the same manner as a normal X.509 end-entity certificate. Moreover, Proxy Certificates usually have a very limited lifetime (generally a few hours) to mitigate the impact of an eventual security breach.

In the EGI Federation [1], the users need to be members of a Virtual Organization (VO) [10] in order to access the resources: a VO is a way of grouping users usually working on the same project and using the same application software. After an agreement with resources providers, VOs have been granted usage of a specific set of resources and services in the infrastructure.

When creating a proxy, an Attribute Authority (AA) can be contacted to release and attach to the proxy the attributes required (if allowed) to access the resources. In the EGI e-Infrastructure, the AA is implemented by VOMS [11]: several VOMS servers, hosting the VOs and information about the enrolled users, are operated by the Resource Centres members of the infrastructure. This kind of proxy carrying the VO attributes is therefore commonly named "VOMS proxy".

2.2 Robot Certificates

Rather than accessing e-infrastructure services directly, users can access them via a portal (or a "Science Gateway"), which can provide a more accessible interface to the services. Quite often user portals provide users with the capability of using institutional credentials to authenticate themselves; then the portal authenticates to the e-infrastructure services by mapping these credentials to the so-called robot certificates [12]. The robot

certificates⁵ are owned by an individual (often the VO manager) who is accountable for the robot operations. In this way, it is not necessary for a user to request a personal X.509 certificate and the registration to a VO, often perceived as a burden due to the bureaucracy: this contributes to increase the user-friendliness of the platforms. Use of robot certificates is then internally accounted for by the portals in compliance to the VO Portal policy.

3 Issues and Challenges for Interoperable AAI

Controlling access to research-related resources and collaborative tools is challenging, particularly when dealing with research communities that can be geographically dispersed across Europe and the globe. The growth of identity federations at the national and international level has proved to be a successful model to efficiently increase scientific collaboration. An identity federation is intended as any number of organizations agreeing to interoperate under a certain rule, a federation policy, set to authenticate and authorise users. Federations are usually circles of trust in which each organisation agrees to trust the Identity Management of the other members.

In this section, we give a quick depiction on the barriers that communities usually face to adopt and use federated access, and what are the common requirements for implementing an interoperable AAI framework.

The AARC project [3], an EC-funded H2020 project, made a survey with 14 European scientific research communities and conducted interviews with a selected but broad representation of user communities in order to understand the most common issues and to classify the several requirements [13].

According to the AARC report [13] and the FIM4R (Federated Identity Management for Research) whitepaper [14], the communities, in general, viewed the Federated Identity Management as an important mean to enable access to shared resources, but the most perceived barrier was the lack of adequate information about it (this highlighted the need to provide guidelines and training, as well as online resources and material for management and decision-makers to facilitate AAI appropriation by each community). Other important barriers were the lack of funding⁶, the excessive bureaucracy when joining a federation, and the lack of clarity on benefits within the organization. Moreover, the survey confirmed that the web-based authentication method cannot solve alone the AAI challenge for VOs: many users still prefer non-web-based authentication, as well as protocol translation and delegation. Most communities reported also that Identity Federations' coverage for their collaboration is poor.

From the interviews and the discussion about the requirements, it was clear that, besides the need to cover functional gaps between the communities, building a federated AAI requires the definition of common policies that cover the necessary legal and operational practices for all the entities involved in the AAI ecosystem. The outcome

⁵ Since the portal is an automated entity, the e-infrastructure services consider it to be a "robot". The portal operator obtains a "robot certificate" that enables the portal to authenticate to e-infrastructure services.

⁶ Often, institutes do not have enough funding for paying the necessary resources and full-time staff to manage them.

of the analysis and prioritization of these requirements was a fundamental input for the high-level AAI Blueprint Architecture, which will be described in the next section.

The requirements were classified into two categories: (A) architectural and technical, and (B) policies and best practices.

In the first category we have:

- **User and Service Provider friendliness:** the Federated AAI framework should provide simple and intuitive tools that are able to address the needs of users with different levels of computer literacy and enable more Service Providers (SPs) (commercial and non-commercial) to connect.
- **Homeless users:** the Federated AAI framework should support users without a federated institutional Identity Provider (IdP), such as citizen scientists and researchers without formal association to research laboratories or universities.
- **Different Levels of Assurance:** credentials issued under different policies and procedures should include the provenance of the level under which they were issued.
- **Community-based authorisation:** the Federated AAI framework should enable communities to manage the assignment of attributes to their members for authorisation purposes.
- **Attribute aggregation/Account linking:** the Federated AAI framework should support the aggregation of identity attributes originating from different sources of authority, including federated IdPs and community-based attribute authorities.
- **Federation solutions based on open and standards-based technologies:** open and standards-based AAI technologies should be used by the different communities to allow for interoperability by means of suitable translation services
- **Persistent user identifiers:** the Federated AAI framework should reference the digital identities of users through long-lasting identifiers.
- **Unique user identities:** Each user should have a single digital identity to allow SPs to uniquely identify their users.
- **User-managed identity information:** A user should be able to self-manage some of their attributes, e.g. through a web-based User Interface (UI). Depending on the attribute type, update restrictions should be imposed.
- **User groups and roles:** the Federated AAI framework should support the assignment of groups to users, as well as the assignment of roles to users within their groups.
- **Step-up authentication:** the Federated AAI framework should provide an additional factor or procedure that validates a user's identity for high-risk transactions or according to policy rules.
- **Browser & non-browser based federated access:** the Federated AAI framework should provide federated access to both web-based and non-web-based services/applications.
- **Delegation:** the Federated AAI framework should provide the capability for the users to delegate to third parties, mostly computational tasks or services, to act on their behalf. This allows users to run thousands of actions in parallel without the need for interactive access, for example, to save output data (as described in Sect. 2.1).
- **Social media identities:** the Federated AAI framework should support common social media providers, such as Google and LinkedIn, but also the researcher identity providers, such as ORCID, to act as authentication providers and/or attribute authorities.

- **Integration with e-Government infrastructures:** the Federated AAI framework should support broader cross-domain collaboration including e-Government infrastructures.
- **Effective accounting:** the Federated AAI framework should support effective accounting across distributed, heterogeneous data infrastructures.

In the second category, policies and best practices, the requirements are the following:

- **Policy harmonisation:** all participating entities in the AAI ecosystem (IdPs, AAs, SPs) should commit to a common policy framework regarding the processing of personal data. This framework should incorporate at least the GÉANT Data Protection Code of Conduct [15].
- **Federated incident report handling:** A common procedure should be adopted for reporting security incidents that involve federations spreading across multiple administrative domains.
- **Sufficient attribute release:** the set of attributes released to SPs should be extended, primarily, to allow consuming services to operate and, also, to allow for more advanced features, such as personalisation of services.
- **Awareness about identity federations:** the benefits offered by identity federations should be promoted to all stakeholders, such as (commercial) service providers and identity providers that have not joined a federation yet.
- **Semantically harmonised identity attributes:** a common set of vocabularies should be used by the different communities to denote identity attributes managed by identity providers and attribute authorities.
- **Simplified process for joining identity federations:** the bureaucracy involved in joining identity federations should be reduced.
- **Best practises for terms and conditions:** AARC could offer guidelines for describing the terms and conditions that service providers (operated in the R&E) should use.

4 A General Solution: The AARC Blueprint Architecture

The way researchers collaborate can vary significantly between different scientific communities. Some are highly structured, with thousands of researchers who could be located virtually anywhere in the world. Typically, these are communities that have been working together for a long time, that want to share and have access to a wide range of resources, and have had to put in place practical solutions to make the collaborations work. On the other hand, there are also a number of smaller, more diverse research communities working within specific or across multiple scientific disciplines. Typically, these are either nascent communities being established around new scientific domains or communities in specific domains that do not need to promote widespread and close collaboration among researchers. In between these two extremes are scientific communities of all varieties in terms of size, structure, history, etc.

Over the past few years, the AARC project [3] has been working together with e-infrastructures, research infrastructures, research communities, AAI architects, and implementers to get a better understanding of their experiences and needs regarding

sharing and accessing resources within research collaborations. The goal has been to collectively define a set of architectural building blocks and implementation patterns, the “AARC Blueprint Architecture” (BPA), that will allow the development of interoperable technical solutions for international intra- and interdisciplinary research collaborations.

Research infrastructures and e-infrastructures can already rely on eduGAIN [2] and the underlying identity federations to authenticate their users: the AARC BPA builds on top of eduGAIN and adds the functionality required to support common use cases within research collaborations, such as access to resources based on community membership.

While previous versions of the BPA [16] provide a blueprint for implementing an AAI, the latest iteration of the BPA (AARC-BPA-2019) [17] focuses on the interoperability aspects, to address an increasing number of use cases from research communities requiring access to federated resources offered by different infrastructure providers. Hence the “community-first” approach, which introduces the Community AAI. The purpose of the Community AAI is to streamline researchers’ access to services, both those provided by their own infrastructure as well as services shared by other infrastructures. User authentication to the Community AAI uses primarily institutional credentials from national identity federations in eduGAIN, but, if permitted by the community, can also use other IdPs.

Specifically, in the community-first approach, we can distinguish among three types of services that can be connected to the Community AAI:

1. community services - provided only to members of a given community
2. generic services - provided to members of different communities
3. infrastructure services - provided by a given research infrastructure or e-Infrastructure to one or more Community AAI (typically through a dedicated infrastructure proxy).

AARC-BPA-2019 [17] is accompanied by a set of guidelines and informational documents that provide guidance on the interoperable expression of information, including

- community user identifiers [18]
- group membership and role information [19]
- resource-specific capabilities [20]
- affiliation information [21].

4.1 The AARC Blueprint Architecture Building Blocks

The current BPA version champions a proxy⁷ service architecture in which services in a research collaboration can connect to a single point, the SP-IdP-Proxy (hereafter termed “proxy”), which itself takes the responsibility for providing the connection to the identity federations in eduGAIN, thus reducing the need for each service having to separately

⁷ Not to confuse with the proxy certificate mentioned in Sect. 2. Hereafter the word “proxy” is meant as “proxy server”: a computer server or an application that acts as an intermediary for requests from clients seeking resources from other servers.

connect to a federation (eduGAIN). As shown in Fig. 1, the latest iteration of the AARC Blueprint Architecture (AARC-BPA-2019) [17] defines five-component layers: User Identity, Access Protocol Translation, Community Attribute Services, Authorisation and End Services. Each layer groups one or more components based on their functional role.

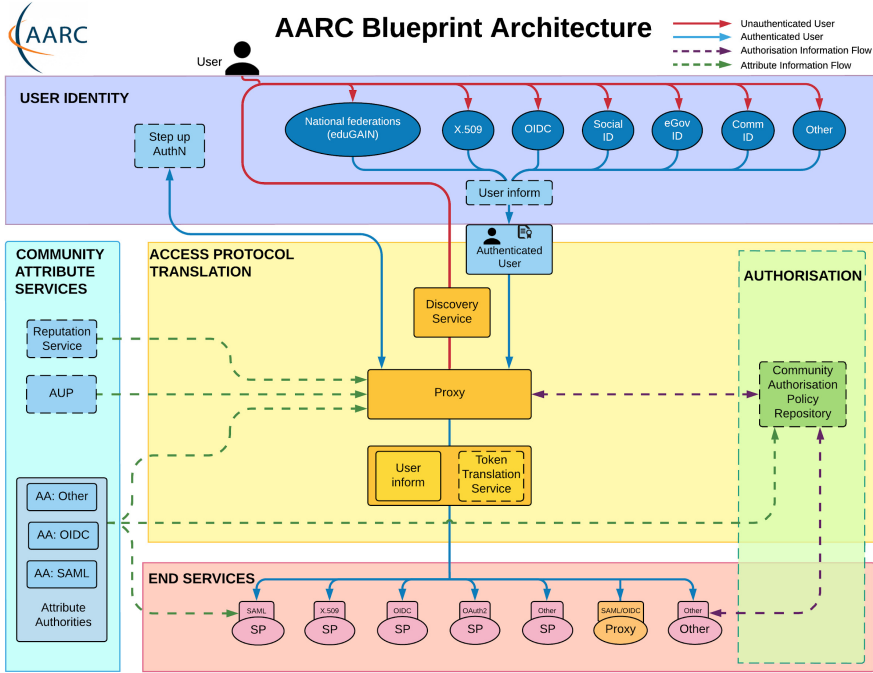


Fig. 1. AARC Blueprint Architecture (AARC-BPA-2019).

The *User Identity Layer* contains services for the identification and authentication of users. In existing implementations in the research and education space, these services typically include Security Assertion Markup Language (SAML) identity providers, certification authorities, and OpenID Connect (OIDC) or OAuth2 Providers (OPs). Although the focus of the services in this layer is to provide user authentication, often some end-user profile information is released as part of the authentication process.

The *Community Attribute Services Layer* groups services related to managing and providing information (attributes) about users. Typically, they provide additional information about the users, such as community group membership and roles, on top of the information that might be provided by services from the User Identity Layer.

The *Access Protocol Translation Layer* addresses the requirement for supporting multiple authentication technologies. It includes the following services:

- SP-IdP-Proxy (proxy), which serves as a single integration point between the Identity Providers from the User Identity Layer and the Service Providers in the End Services

Layer. Thus, the proxy acts as an SP towards the Identity Federations for which this proxy looks like any other SP, while towards the internal SPs it acts as an IdP.

- Token Translation Services, which translate identity tokens between different technologies.
- Discovery Service, which enables the selection of the user's authenticating IdP.
- User inform, which allows users to be informed regarding the processing of their personal data.

The *Authorisation Layer* controls access to the End Services Layer. The AARC BPA allows the implementers to delegate many of the complex authorisation decisions to central components, which can significantly reduce the complexity of managing authorisation policies, and their evaluation for each service individually.

The *End Services Layer* contains the services users want to use. Access to these services is protected (using different technologies). These services can range from simple web-browser-based services, such as wikis or portals for accessing computing and storage resources, to non-web-browser-based resources such as APIs, login shells, or workload management systems.

4.2 The “Community-First” Approach

As mentioned above, the latest BPA iteration fosters the interoperability among AARC BPA compliant AAI that are operated by different research and e-Infrastructures by introducing the so-called Community AAI, which follows the proxy-based architecture shown in Fig. 1. It is therefore responsible for dealing with the complexity of using different identity providers with the *community services*. Furthermore, the Community AAI can add attributes to the federated identity that in turn can enable service providers to control access to their resources. These community-specific services only need to connect to a single identity provider, i.e. their Community AAI.

Apart from the community-specific services, there are *generic services*, such as the RCauth.eu Online CA, which serve the needs of several communities and are thus connected to more than one Community AAI. Being connected to multiple Community AAI requires generic services to provide some form of IdP discovery, in order to be able to redirect the user to the relevant Community AAI⁸. Additionally, the generic services should support some means of doing “IdP hinting” (see [22]), thereby allowing “community branding” of the service and automatically redirecting the user to the corresponding Community AAI.

Communities may also require access to various services which themselves are behind (another) proxy, as often is the case with resources offered by e-Infrastructures or Research Infrastructures (Infrastructures hereafter). These *Infrastructure Proxies*⁹

⁸ Primarily to get the user's identity via the community IdP, but also potentially to obtain attributes from community attribute authorities.

⁹ An AAI service of a research infrastructure or e-Infrastructure (hereafter termed infrastructure) that enables access to resources offered by Service Providers connected to that infrastructure. This AAI service does not provide community membership management. Specifically, the infrastructure proxy comprises two AARC BPA component layers: the Access Protocol Translation and the Authorisation.

can be connected to different Community AAIs - see Fig. 2. So, just as for the generic services, Infrastructure services should be able to hint to the Infrastructure Proxy which Community AAI to use (see [22]).

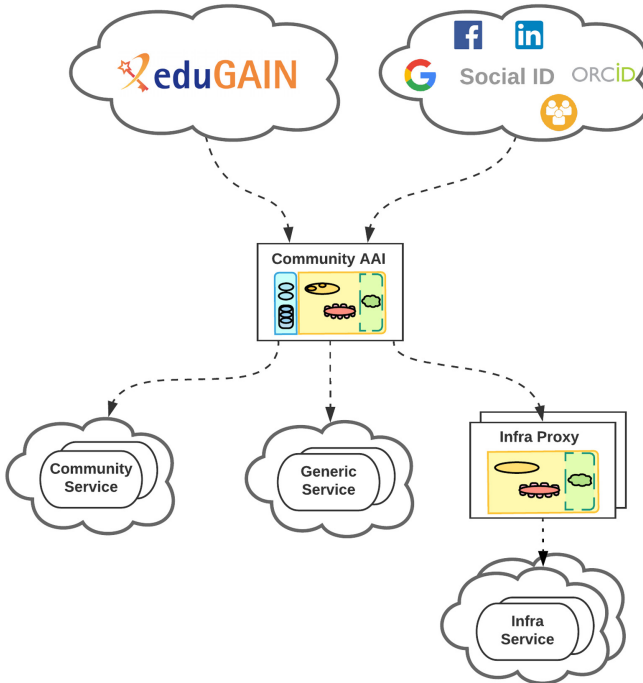


Fig. 2. A Community-first approach based on the AARC Blueprint Architecture. Researchers access services/resources using their institutional (eduGAIN), social or community-managed IdP via their Community AAI. Community services are connected to a single Community AAI, whereas generic services can be connected to more than one Community AAIs. e-Infrastructure services can be connected to different Community AAIs through a single e-infrastructure SP proxy. (A community-managed IdP is useful when there is a collaboration that wants to release attributes at IdP level for its members. This would allow to streamline the authentication process at community level. It can also be useful when a large part of the collaboration members does not have their own identity provider.)

It should be noted that the “community-first” approach does not impose a requirement on communities to deploy and operate a Community AAI on their own. Communities could make use of either dedicated or multi-tenant deployments of AAI services operated by a third-party, typically a generic e-Infrastructure. A multi-tenant AAI service deployment supports different communities, as depicted in Fig. 3. It typically appears as a single entity to its connected IdPs and SPs. Such multi-tenant deployments are aimed at medium-to-small research communities/groups or individual researchers. Yet it should be emphasised that also in the multi-tenant AAI scenario, the community managers are responsible for managing their community members, groups and authorisation attributes.

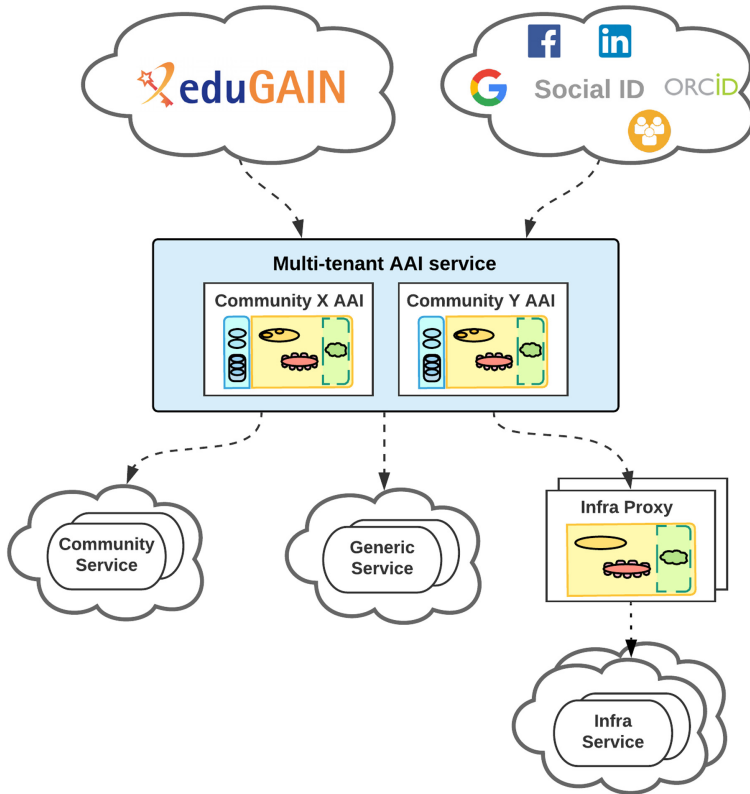


Fig. 3. Multi-tenant deployment of AAI services in “community-first” approach to the AARC Blueprint Architecture.

4.3 Authorisation Models

Authorisation models describe the organisational flow of authorisation information. Any other information needed by the service to fulfil actions such as personalisation, accounting, traceability, is out of the scope of this chapter. The organisational flow of authorisation information follows this lifecycle:

- Definition of authorisation information at one or more Attribute Authorities (AA)
- Aggregation of authorisation information
- Use of authorisation information for making an authorisation decision
- Enforcement of the authorisation decision

Authorisation information can be classified into two types:

1. User-attributes (often aggregated from different sources) such as:
 - Affiliation within the Home Organisation and/or the Community
 - Assurance, i.e. how well attribute assertions can be trusted
 - Group and role information (these primarily come from the Community)
2. Capabilities such as information describing what actions a user is entitled to perform on a specific resource.

Based on the analysis of the authorisation architectures from nine different use cases detailed in [23], it has been identified three main authorisation models that make use of an SP-IdP-Proxy, as shown in Fig. 4:

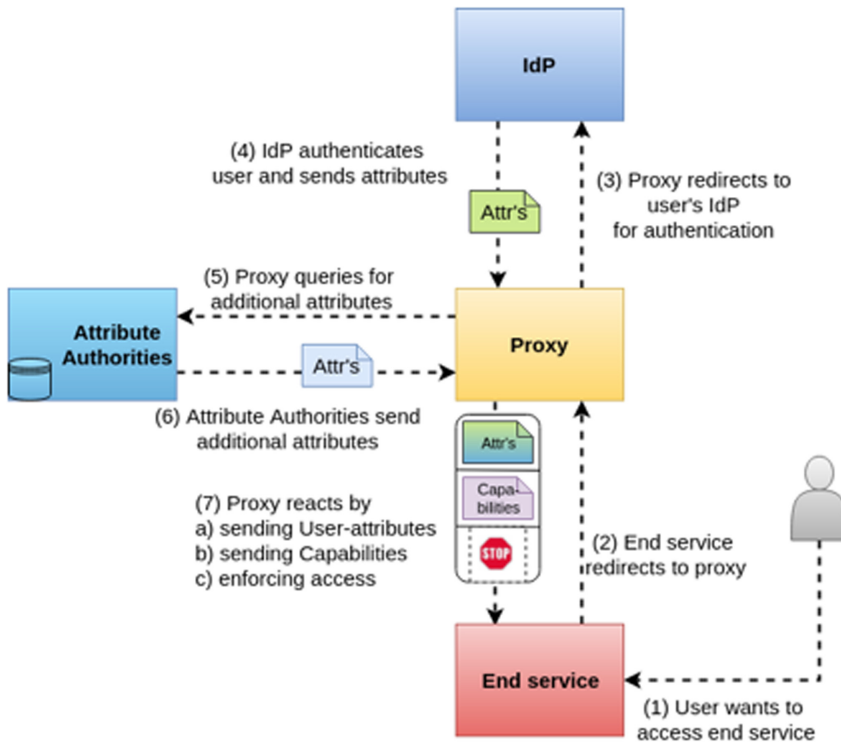


Fig. 4. The flow of authorisation information for a user who wants to access an end service in a BPA-compliant infrastructure.

1. Centralised Policy Information Point (step 7a in Fig. 4): the proxy aggregates user attributes, such as group membership information and roles, and makes them available to the end-services

2. Centralised Policy Management and Decision Making (step 7b in Fig. 4): the proxy conveys the authorisation decision to the end-services in the form of capabilities
3. Centralised Policy Management and Decision Making and Enforcement (step 7c in Fig. 4): the proxy enforces the decision directly at the proxy.

Centralised Policy Information Point. In this model, the proxy aggregates the information and makes it available to the end services so they can make the authorisation decision. This allows the service to perform fine-grained access control because all information necessary for an informed decision is available. However, scalability may become an issue for large deployments. For example, it may become non-trivial to consistently update authorisation across a large number of services, as the authorisation policy needs to be replicated to every service. Additionally, services may see user-specific authorisation data, such as group membership, that might be intended for other services. This may be problematic with regard to the “data minimisation principle”. Furthermore, this puts the onus on the services to correctly interpret and act on the obtained authorisation information.

Centralised Policy Management and Decision Making. In this model, the proxy makes the authorisation decision and encodes this decision into resource-specific authorisation information, typically in the form of capabilities. This allows the decision at the proxy to be based on additional information which the proxy might prefer not to send to the services. This is generally simpler for the end services to implement since the complexity of interpretation of the authorisation information is handled by the proxy. In contrast to the approach described in the previous model, this puts the onus on the proxy to correctly interpret and act on the authorisation information. Note that in this model:

1. the proxy is creating and/or translating authorisation statements
2. the proxy may need to make a mix of capabilities and user attributes available for the service to be able to properly enforce the authorisation decision.

Centralised Policy Management and Decision Making and Enforcement. In this model, the proxy makes the authorisation decision, as in the case of Centralised Policy Management and Decision Making. Furthermore, the proxy is responsible for enforcing that decision. This allows the integration of services that might not be capable of doing any authorisation, with only little modification. However, it requires the proxy to understand the authorisation policy of the end services. Often this type of authorisation enforcement is only used for certain parts (e.g. a global black- or whitelist) while using the other models for the rest of the authorisation. For example, in case the proxy grants the user access to the end service, this model may be followed by either of the other two models described.

Considerations on the Different Models

1. Authorisation implementations **SHOULD** support the Centralised Policy Information Point model for end services that require full control over the authorisation process. Authorisation implementations **MUST** be aware that in this model it is easy to send more data than required to end service. Filtering **MAY** be a solution.

2. Authorisation implementations SHOULD support the Centralised Policy Management and Decision Making model for simplifying the authorisation process for the end services. Authorisation implementations MUST be aware that the onus for correctly interpreting and acting upon authorisation information is put on the proxy.
3. Authorisation implementations SHOULD only use the Centralised Policy Management, Decision Making and Enforcement model for a partial authorisation decision (e.g. central suspension), and combine it with one of the two models above.
4. Depending on the requirements of the Service Providers reached through the proxy, it is possible to use a hybrid approach, combining any of the three models above, in a single authorisation flow. In all these flows the proxy can supplement the attributes from the authenticating IdP with information from AAs. The three different approaches address whether and how this information is passed on to the end services.

5 The EGI AAI Platform

The Check-in service is the AAI Platform for the EGI infrastructure [24] that implemented the AARC Blueprint Architecture. The Check-in service enables the integration of external Identity Providers (e.g. from eduGAIN [2] and individual organisations) with the EGI services through the Check-in Identity/Service Provider Proxy component, so that users are able to access the EGI services (web and non-web based) using existing credentials from their home organisations. To this end, Check-in has been published in eduGAIN as a Service Provider. Through eduGAIN, EGI operational tools and services that are connected to Check-in can become available to more than 3000 Universities and Institutes from the 60 eduGAIN Federations with little or no administrative involvement.

Compliance with the REFEDS Research and Scholarship (R&S¹⁰) entity category and the Sirtfi¹¹ framework, the Check-in service ensures sufficient attribute release, as well as operational security, incident response, and traceability for 170 Identity Providers from 25 identity federations that support R&S and Sirtfi. Complementary to this, users without an account on a federated Identity Provider are still able to use social media or other external authentication providers for accessing EGI Services that do not require substantial level of assurance [25].

The adoption of standards and open technologies by Check-in, including SAML 2.0¹², OpenID Connect¹³ and X.509 v3, has facilitated interoperability and integration with the existing AAIs of other eInfrastructures and research communities, such as

¹⁰ The REFEDS Research and Scholarship Entity Category (R&S) is one of the Entity Categories defined by REFEDS, <https://refeds.org/category/research-and-scholarship>.

¹¹ Sirtfi - A Security Incident Response Trust Framework for Federated Identity, defined by REFEDS <https://refeds.org/wp-content/uploads/2016/01/Sirtfi-1.0.pdf>.

¹² SAML 2.0 standard is produced by the SSTC on 1 May 2012: https://wiki.oasis-open.org/security/FrontPage#SAML_V2.0_Standard.

¹³ OpenID Connect 1.0 is a simple identity layer on top of the OAuth 2.0 protocol. It allows Client to verify the identity of the End-User based on the authentication performed by an Authorisation Service. The OpenID specification is at <http://openid.net/developers/specs/>.

ELIXIR¹⁴ and LToS¹⁵. The Check-in service enables users to manage their accounts from a single interface, to link multiple accounts/identities together and to access the EGI services based on their roles and Virtual Organisation (VO) membership rights. For VOs that do not operate their own Group/VO management system, the Check-in service provides an intuitive interface to manage their users and their respective roles and group rights. For VOs that operate their own Group/VO management system, the Check-in service has a comprehensive list of connectors that allows integrating their systems as externally managed Attribute Authorities (AA).

In summary, user communities have several options to integrate with Check-in in order to access the EGI resources:

- Users authenticate using their institutional identity provider, which is part of an identity federation and eduGAIN;
- Users authenticate using their ORCID, social media a community-specific identity provider, for example, in the case of ELIXIR;
- Users authenticate using their Community AAI (see also Sect. 4.2), for example, in the case of ELIXIR;
- Authorisation information about the users (VO/group memberships and roles) is managed by the community's group management service, which is connected to Check-in as an external attribute authority;
- Communities that do not operate their own Group/VO management service can leverage the group management capabilities of the Check-in platform.

EGI Check-in is a contribution towards the development of Single Sign On (SSO) to e-infrastructures for European researchers. It lowers the barriers to use of EGI resources today and has been designed with an eye to the integration with other planned and probable developments. Check-in service can be accessed at <https://aai.egi.eu/>.

5.1 EGI Check-in Architecture

Figure 5 illustrates a high-level view of the Check-in architectural elements that deliver the system's functionality. It depicts the system's functional structure, including the key functional components, their responsibilities, the interfaces they expose, and the interactions between them.

The core of EGI AAI Check-in service is the **IdP/SP Proxy** component, which acts as a bridge between the EGI services and external authentication sources and identity providers. This decoupling of the internal services and the external authentication

¹⁴ ELIXIR: A Europe leading life science organisations in managing and safeguarding the data being generated by publicly funded research. <https://www.elixir-europe.org/>.

¹⁵ LToS: The long-tail of science refers to the individual researchers and small laboratories who - opposed to large, expensive collaborations - do not have access to computational resources and online services to manage and analyse large amounts of data. EGI provides the Application on Demand (AoD) service, which is a platform allows individual researchers and small research teams to perform compute and data-intensive simulations on large, distributed networks of computers in a user-friendly way. https://wiki.egi.eu/wiki/Long-tail_of_science.

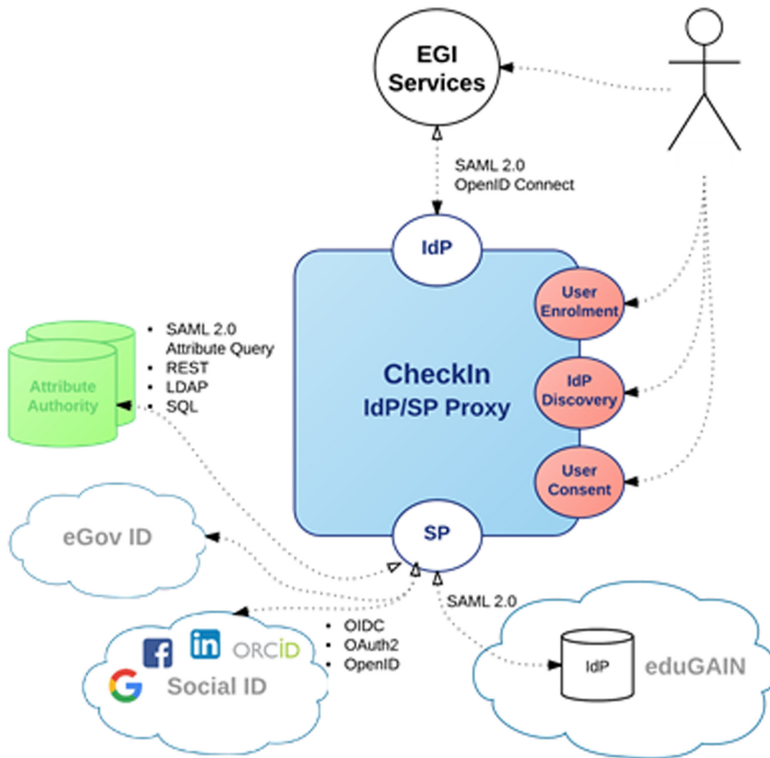


Fig. 5. EGI checkIn high-level functional architecture.

sources/identity providers reduce the complexity of the service implementation as it removes dependencies on the heterogeneity of multiple IdPs, Federations, Attributes, Authorities and different authentication and authorization technologies. This complexity is handled centrally by the proxy.

The introduction of an IdP/SP Proxy entity brings additional benefits. Specifically, as illustrated in Fig. 5, services only need to establish trust with one entity, the IdP/SP proxy. Typically, services will have one static configuration for the IdP/SP proxy. Having one configured IdP also removes the requirement from the service providers to operate their own IdP Discovery Service (a common requirement for services supporting federated access). Furthermore, all internal services will get consistent and harmonised user identifiers and attributes, regardless of the home organisation or the research community the authenticating user belongs to. Finally, this separation simplifies change management processes, as the internal services are independent of the IdPs run by the home organisations. Similarly, IdPs establish trust with one entity, the operator of the IdP/SP proxy, and they are not impacted by the operational changes introduced by each individual service.

The **User Enrolment and VO Management** service supports the management of the full life cycle of user accounts in the Check-in service. This includes the initial user registration, the acceptance of the terms of use of EGI, account linking, group and VO

management, delegation of administration of VOs/Groups to authorised users and the configuration of custom enrolment flows for VOs/Groups via an intuitive web interface.

5.2 Token Translation: Integration with RAuth.Eu Online CA

For various use cases, a user might need to use different types of credentials: for example, the user has an institutional account but she needs to access a storage element that requires an X.509 (proxy) certificate. So it is necessary to translate those institutional credentials into the precise format allowing the access to that particular service. In order to provide such functionality, the EGI Check-in service has been connected to the new RAuth.eu Online CA [26].

The RAuth Online CA issues certificates to end-entities based on a successful authentication to a Federated Identity Management System (FIMS) operated by an eligible Registration Authority – typically a FIMS Identity Provider (IdP) operated by an academic or research organisation.

When a certain web-flow requires a X.509 credential, the user will be redirected via a component, a so-called Master Portal¹⁶, to the Online CA¹⁷. There the user will log in again transparently (due to SSO) to the Check-in service and will have to give consent for the management of user credentials. It will then be redirected to the originating service. In the process, a new credential is cached in the Master Portal which subsequently will be retrieved by whichever service initiated this flow, typically a Science Gateway.

When it is needed, a VOMS proxy (as seen in Sect. 2.1) can be requested initially. When the user is already enrolled in the VOMS server, this can be done completely transparently; otherwise a form of provisioning is needed.

The components of the service, as shown in Fig. 6, can be categorised in the following way:

- The **blue component** [many] represents the Service Provider Portal which the user wants to use. These are usually the Science Gateways (VO Portals) run by VOs. Given the wide variety of scientific disciplines in EGI, this scenario may include many such portals.
- **Red components** [few] correspond to the Master Portal. The scenario may comprise a few of these services, each one corresponding to the e-Infrastructures (like EGI) using RAuth.eu.
 - *Master Portal*: acts as a caching service for user credentials (proxy certificates), taking some load of the RAuth.eu backend. Moreover, it also intermediates between two separate trust domains: the domain (single) of the Delegation Server and the domains (many) of connecting Service Provider Portals (Science Gateways). This improves the scalability of the model since instead of registering ALL Portals to the single Delegation Server directly, now registered Portals can be split between a few Master Portals running in front of the Delegation Server.

¹⁶ The architecture design of the Master Portal is at: https://wiki.nikhef.nl/grid/AARC_Pilot_-_Architecture.

¹⁷ The architecture of the RAuth Online CA is at: https://wiki.nikhef.nl/grid/AARC_Pilot_-_RAuth.eu.

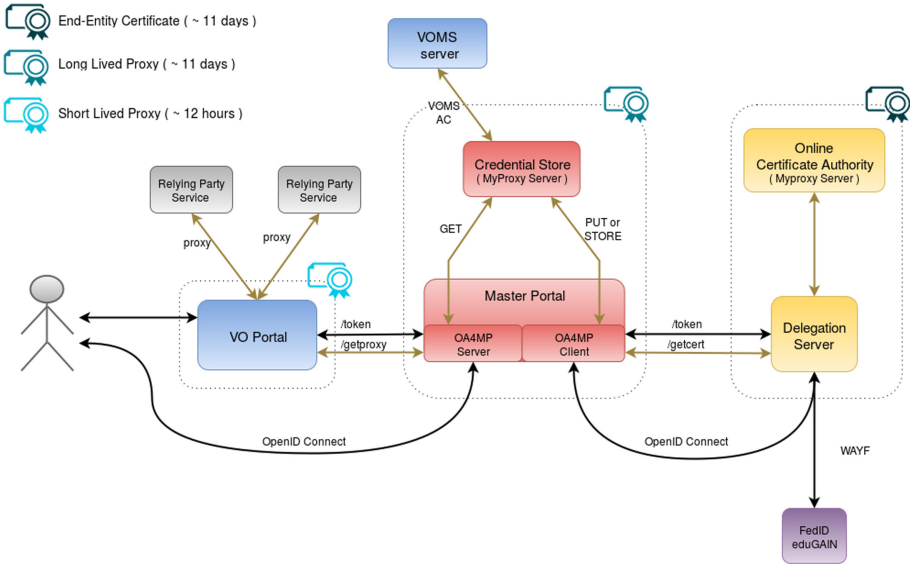


Fig. 6. RAuth.eu Online CA scenario [27]. (Color figure online)

- *Credential Store*: is a MyProxy server used by the Master Portal to actually store the user proxies.
- **Yellow components** [one] represent an Online CA with a web frontend, and it's what we call RAuth.eu. Given the hardware security module (HSM) cost and high-security requirement, there is only one Online CA component.
 - *Delegation Server*: is the web frontend service which talks to the Online CA to generate certificates for authenticated users.
 - *Online CA*: is a Certificate Authority running on an HSM. This service (although called online) is only directly accessible to the Delegation Server in front of it.
 - *WAYF*: an IdP/ SP Proxy with an internal filter for accepting authentication sources directly. This gives full control for RAuth.eu over the eligibility of IdPs.
- **Purple components** [many] represent the different authentication sources that RAuth.eu is accepting (or planning to accept in the near future).

The addition of the Master Portal component to the schema, based on a replication of the CILogon software¹⁸, moves all the complexity of caching the user credentials and of interacting with the Online CA away from the VO-run science gateways. The net result is that it makes easy for the VO portals to securely obtain credentials, based on the OpenID Connect protocol (see footnote 9) (acting as a client). The Master Portal takes

¹⁸ CILogon is an integrated identity and access management platform that enables researchers to log on to cyberinfrastructure (CI): <https://www.cilogon.org/>.

care of obtaining the longer-lived end-entity certificates, caching them in the form of a proxy certificate and handling the additions of the VO-based attributes. Due to the more modular setup, having this extra component in the middle also makes it easier to reuse the same online CA for different e-Infrastructures.

6 Accounting

In previous sections, we have discussed two aspects of AAA. The final plank in the AAA framework is accounting, which measures the resources a user consumes during access. This can include the amount of system time or the amount of data a user has sent and/or received during a session. Accounting is carried out by logging of session statistics and usage information and is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities [28]. Accounting is fundamental in measuring the resource usage for each VO and verifying it's in line with the SLAs and the corresponding requirements/pledges negotiated with it. Moreover, EGI's "pay for use" model, which supposes that resources are paid by the customer periodically as they are consumed, will surely make use of accounting data in the future, as soon as the volumes of usage will be high enough.

For the purpose of this book, we are going to provide an overview of the accounting implementation in EGI.

The EGI Accounting Infrastructure (Portal and Repository) supports the daily operations of EGI and it is useful for assessing the real usage of the computing, cloud, and storage resources.

It is a complex system that involves various sensors in different regions, all publishing data to a central repository. The data are processed, summarised and displayed in the accounting portal, which acts as a common interface to the different accounting record providers and presents a homogeneous view of the data gathered and a user-friendly access. There are dedicated views for different types of users, for example national resources managers, Virtual Organisation (VO) Managers, resource centres administrators and the general public.

The Accounting Repository is based on APEL [29], a tool that collects accounting data from sites participating in the EGI. The accounting information is gathered from different sensors into a central accounting database where it is processed to generate statistical summaries that are available through the EGI Accounting Portal¹⁹ [4]. Statistics are available for view in different detail by users, VO Managers, site administrators and anonymous users according to well-defined access rights.

The Accounting Portal is a web application based on Apache, and MySQL, which has as its primary function to provide users with customised accounting reports, containing tables and graphs, as web pages. It also offers RESTful web services to allow external entities to gather accounting data.

¹⁹ EGI Accounting Portal is one of EGI core services that provide data accounting information for EGI users: <https://accounting.egi.eu/>.

The Accounting Portal consists of a backend (Fig. 7), which aggregates both data and metadata in a MySQL database, using the APEL SSM (Secure Stomp Messenger)²⁰ messaging system to interact with the Accounting Repository and several scripts, which periodically gather the data and metadata. It relies on a model that allows the representation of the data in several ways, focusing on different views (grid, cloud, storage, multicore, user statistics etc.) and integrating metadata (topology, geographical data, site status, nodes, VO users and admins, site admins etc.). Secure Stomp Messenger (SSM) is based on Apache ActiveMQ²¹.

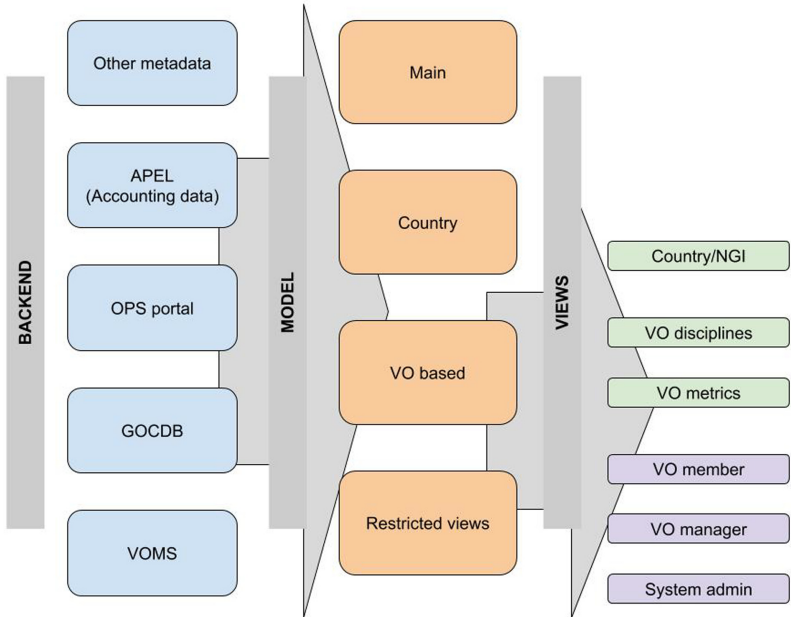


Fig. 7. Accounting Portal information sources and the different views provided.

A set of specific views exposes the data to the user. These views contain a form to set the parameters and metric of the report, a number of tables showing the data parametrised by two selectable dimensions and filtered by several parameters, a line graph showing the table data, and pie charts showing the percentage distribution on each dimension.

The Accounting Portal has to refresh its database periodically with data from the Accounting Repository to ensure that information published are up-to-date.

Metadata is a category of data that complements the raw accounting data and allows the portal to organise, categorise and import new meaning to it. This metadata includes:

²⁰ APEL SSM is the messaging system used by APEL to transmit messages: <https://wiki.egi.eu/wiki/APEL/SSM>.

²¹ Apache ActiveMQ is a multi-protocol, java-based messaging server, <http://activemq.apache.org/>.

- **Geographical Metadata:** Country and Operations Centre affiliation of sites. Generally, this follows current borders, but there are important exceptions.
- **Topological Metadata:** Sites are presented in trees, there are Country and Operations Centres trees that correspond to geographical classifications.
- **Role Metadata:** VO members and managers, and the site admins records. This metadata controls the access to restricted views.
- **Country affiliation data:** Each user record contains a user identifier that has his/her user name and membership data. These data are used in anonymised statistics per country, like how much resources from other countries are used by a given country and the distribution of its resources used by other countries.
- **VO Data:** To make possible VO selection in the user interface, the portal stores the list of all the VOs. They are also used to filter incorrect VO names, provide access to VO managers, and arrange to account by VO discipline (such as “High Energy Physics”, “Biomedicine”, “Earth Sciences”, etc.). Information is gathered from the Operations portal using its XML based APIs.
- **Site status metadata:** Sites must be filtered to exclude those that are not in production (due to being closed or being in test mode). There must be also metadata to aggregate the accounting history of sites whose name has been changed.
- **Other metadata:** There are also other metadata like local privileges, SpecInt calculations, publication status, VO activities and more. Some of these metadata is calculated internally using other types of metadata and published for other EGI operational tools, like VO activity data.

Views in the portal differ in the type of showed accounting data, the site organization or the restricted nature of data. The Cloud view is a view of the sites that are part of EGI Federated cloud platform, which uses Cloud middleware. Some relevant views to be considered are:

- **The main grid/cloud view,** showing metrics like “Sum elapsed CPU time” or the “Number of jobs”. An important metric to evaluate the performance of a cloud Resource Centre is the “Elapsed time * Number of Processors (hours)”, together with the “Number of VMs” running at that RC.
- **Operations centres and Countries view:** similar to the main view, but showing data per country or per “Operations Centre”: the generic Operations Centre, in general, is mapped to a country, but there are cases where it’s a group of countries (e.g. NGI_IBERGRID), or a fraction of a country (e.g. CERN).
- **Disciplines View:** A view that provides accounting data per VO scientific disciplines defined by EGI.

7 Conclusion

In this chapter, we started with a review of the advanced AAI technology and discussed them with the best practices in the EGI e-Infrastructure. We also brought up the interoperable AAIs and the identity federation issues that challenging today’s science collaborations. We presented the AARC Blueprint Architecture as one of the sound

solutions and provided an implementation example of EGI Check-in service. We finally addressed the last ‘A’ of AAA – Accounting, and described the technology and services used by EGI.

AAAI solutions have been rarely implemented in the ENVRI Research Infrastructures. The experience described here are generic AAA solutions and have been implemented and used by EGI e-Infrastructure to support daily operations. These solutions can be easily extended and adopted by ENVRI RIs.

Acknowledgements. This work was supported by the European Union’s Horizon 2020 research and innovation programme via the ENVRIplus project under grant agreement No 654182.

The work presented in this paper was also supported by the EGI-Engage H2020 project (grant no. 654142), the EOSC-hub project (grant no. 777536), the AARC project (grant no. 653965) in particular for Sect. 3, and the AARC2 project (grant number 730941) in particular for Sect. 4.

References

1. EGI e-Infrastructure Homepage. www.egi.eu. Accessed 30 Apr 2019
2. eduGAIN Homepage. <https://edugain.org/>. Accessed 30 Apr 2019
3. The AARC project Homepage. <https://aarc-project.eu>. Accessed 30 Apr 2019
4. EGI Accounting Portal. <https://accounting.egi.eu/>. Accessed 30 Apr 2019
5. Adams, C., Lloyd, S.: Understanding PKI: Concepts, Standards, and Deployment Considerations, pp. 11–15. Addison-Wesley Professional (2003). ISBN 978-0-672-32391-1
6. The Interoperable Grid Trust Federation (IGTF). <https://www.igtf.net/>. Accessed 30 Apr 2019
7. X.509 Version 3. <https://tools.ietf.org/html/rfc5280>. Accessed 30 Apr 2019
8. The Transport Layer Security (TLS) Protocol Version 1.3. <https://tools.ietf.org/html/rfc8446>. Accessed 30 Apr 2019
9. Proxy Certificate Profile. <https://tools.ietf.org/html/rfc3820>. Accessed 30 Apr 2019
10. Virtual Organization (VO). https://wiki.egi.eu/wiki/Glossary_V3#Virtual_Organisation. Accessed 30 Apr 2019
11. VOMS. <http://italiangrid.github.io/voms/documentation.html>. Accessed 30 Apr 2019
12. Robot Certificates. <https://www.eugridpma.org/guidelines/robot/>. Accessed 30 Apr 2019
13. Kanellopoulos, C., Liampotis, N., van Dijk, N., Solagna, P.: Analysis of user community and service provider requirements. The AARC project Deliverable DJRA1.1 (2015). <https://aarc-project.eu/wp-content/uploads/2015/10/AARC-DJRA1.1.pdf>
14. Atherton, C.J., et al.: Federated Identity Management for Research Collaborations (Version 2.0) (2018). <http://doi.org/10.5281/zenodo.1307551>
15. The GÉANT Data Protection Code of Conduct. <https://www.geant.org/uri/Pages/dataprotection-code-of-conduct.aspx>. Accessed 30 Apr 2019
16. Kanellopoulos, C., Stevanovic, U., Hardt, M., the rest of the JRA1 Team: AARC Blueprint Architectures. The AARC project Deliverable DJRA1.2 (2017). <https://aarc-project.eu/wp-content/uploads/2017/05/DJRA1.2-AARC-Blueprint-Architectures-1.pdf>
17. AARC Consortium Partners, AppInt Members, Liampotis, N. (ed.): Evolution of the AARC Blueprint Architecture. AARC2 project Deliverable DJRA1.4 (2019). https://aarc-project.eu/wp-content/uploads/2019/05/AARC2-DJRA1.4_v2-FINAL.pdf
18. AARC-G025: Exchange of affiliation information between infrastructure. <https://aarc-project.eu/guidelines/aarc-g025/>. Accessed 30 Apr 2019
19. AARC-G002: Expressing group membership and role information. <https://aarc-project.eu/guidelines/aarc-g002/>. Accessed 30 Apr 2019

20. AARC Consortium Partners, AppInt Members: AARC-G027: Guidelines for expressing resource capabilities (2018). <https://doi.org/10.5281/zenodo.2247446>
21. AARC-G025: Exchange of affiliation information between infrastructures. <https://aarc-project.eu/guidelines/aarc-g025/>. Accessed 30 Apr 2019
22. AARC-G049: A specification for IdP hinting. <https://aarc-project.eu/guidelines/aarc-g025/>. Accessed 30 Apr 2019
23. Hardt, M., et al.: AARC2-DJRA1.2: scalable, integrated authorisation models for SPs. The AARC2 project Deliverable DJRA1.2 (2018). https://aarc-project.eu/wp-content/uploads/2018/07/AARC2-DJRA1.2_V4-FINAL.pdf
24. Kanellopoulos, C., Liampotis, N., Solagna, P., Salle, M.: Identity Management for Distributed User Communities. The EGI-Engage project Deliverable, D3.9 (2017). <https://documents.egi.eu/public/ShowDocument?docid=3017>
25. Groep, D., Jensen, J., Linden, M., Stevanovic, U., Vaghetti, D.: Expression of REFEDS RAF assurance components for identities derived from social media accounts. The AARC2 project Deliverable, AARC_G041 (2018). <https://aarc-project.eu/wp-content/uploads/2018/03/AARC-G041-Expression-of-REFEDS-RAF-assurance-components-for-social-media-accounts.pdf>
26. RCauth.eu Homepage. <https://rcauth.eu/>. Accessed 30 Apr 2019
27. The AARC RCauth.eu pilot Homepage. https://wiki.nikhef.nl/grid/AARC_Pilot_-_RCAuth.eu. Accessed 30 Apr 2019
28. Ren, J., Tongtong, L.: Enterprise Security Architecture, Handbook of Technology Management. Wiley, Hoboken (2010)
29. EGI APEL WiKi page: <https://wiki.egi.eu/wiki/APEL>. Accessed 30 Apr 2019

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

