European Commission

EOSC Authentication and Authorization Infrastructure

Report from the EOSC Executive Board Working Group (WG) Architecture AAI Task Force (TF)

Research and Innovation

Independent Expert Report

EOSC Executive Board
WG Architecture
January 2021

EOSC Authentication and Authorization Infrastructure

European Commission
Directorate-General for Research and Innovation
Directorate G — Research and Innovation Outreach
Unit G.4 — Open Science
Contact Corina Pascu
Email Corina.PASCU@ec.europa.eu
RTD-EOSC@ec.europa.eu
RTD-PUBLICATIONS@ec.europa.eu
European Commission
B-1049 Brussels

Manuscript completed in January 2021.

The European Commission is not liable for any consequence stemming from the reuse of this publication.
The views expressed in this publication are the sole responsibility of the author and do not necessarily reflect the views of the European Commission.

More information on the European Union is available on the internet (http://europa.eu).

PDF	ISBN 978-92-76-28113-9	doi: 10.2777/8702	KI-04-20-744-EN-N

Luxembourg: Publications Office of the European Union, 2021

Cover page: © Lonely #46246900, ag visuell #16440826, Sean Gladwell #6018533, LwRedStorm #3348265, 2011; kras99 #43746830, 2012. Source: Fotolia.com.

EUROPEAN COMMISSION

EOSC Authentication and Authorization Infrastructure

Report from the EOSC Executive Board Working Group (WG) Architecture PID Task Force (TF)

Edited by: the EOSC Executive Board

January 2021

Authors

Klaas Wierenga, Leif Johansson, Christos Kanellopoulos, David Groep, Davide Vaghetti, Nicolas Liampotis

Contributors

Mark van de Sanden, Johannes Reetz, Sadaf Alam, Mikael Linden, Jens Jensen, Niels van Dijk, Andrea Ceccanti, Sander Apweiler, Bjorn Abt, Alex Vermeulen, Patrick Fuhrmann, Marcus Hardt, TF-members

2021 Directorate-General for Research and Innovation

Contents

1 Introduction

The EOSC Architecture Working Group has assigned the AAI Task Force (AAI TF) the task to establish a common global ecosystem for identity and access control infrastructures for the European Open Science Cloud (EOSC).

Since the EOSC is part of an international environment of research and education, the principles established by the EOSC AAI subtask must be globally viable. The EOSC AAI TF has produced a set of deliverables:

- EOSC AAI First Principles & Requirements
- EOSC AAI Baseline Architecture
- EOSC AAI Federation participation guidelines (participation policy and technical framework)
- EOSC AAI Best Practises

This document encompasses the first 3 deliverables, the Best Practises document is a live document that continues to be updated and that can be found at: https://docs.google.com/document/d/1iR56Xb2DBG5JLBLKmva_Si7tcfjjOy__cCFlGRlF0jU

1.1 Governance of this document

This document is the final deliverable of the EOSC AAI Task Force, operating in the context of the EOSC Architecture Working Group[1] and will be handed over to the EOSC Association.

It is recommended that this work will be continued by the EOSC Federation Operator in collaboration with the AARC Engagement Group for Infrastructures (AEGIS)[2] and the pertinent EOSC Task Forces addressing its authentication and authorization infrastructure.

[1] https://www.eoscsecretariat.eu/working-groups/architecture-working-group

[2] https://aarc-project.eu/about/aegis/

2 DEFINITIONS AND TERMINOLOGY

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119[3]].

Common terms and definitions used in the context of the EOSC AAI:

AAI service: A service that enables authenticated and authorised access to resources.

Community: A group of users, organised with a common purpose, and jointly granted access to resources. It may act as the interface between individual users and the resources. (see also [WISE-SCI])

Community AAI: An AAI service that also enables the use and management of community identities for access to resources. It comprises three (3) AARC BPA component layers: the Access Protocol Translation, the Community Attributes Services, and the Authorisation.

Community identity: A user's digital identity that may be enriched by the community with additional attributes such as a shared user identifier, profile information, and community attributes such as group membership and role information (see [REFEDS-R&S] and [REFEDS-Sirtfi]).

Community service: A service provided to members of a specific community.

Digital identity: Information that represents an entity (subject) within a domain. It contains information about the subject's attributes and relationships.

Entity: A discrete AAI component that can be added in the EOSC AAI Federation Registry

Federation: Identity Federation. An association of organisations that come together to securely exchange information as appropriate about their users and resources to enable collaborations and transactions.

Federation member: An organisation that has joined the Federation by agreeing to be bound by the Federation Policy.

Federation Operator: Organisation providing the infrastructure for Authentication and Authorisation to Federation Members.

Federation policy: A document describing the obligations, rights and expectations of the federation members and the federation Operator.

Generic service: A service provided to users, possibly as members of different communities.

Infrastructure service: A service provided by a research infrastructure or e-Infrastructure to members of one or more Community AAI which receives the required attributes through an Infrastructure Proxy.

Infrastructure Proxy: 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

[3] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

membership management[4]. Specifically, the Infrastructure Proxy comprises two (2) AARC BPA component layers: the Access Protocol Translation and the Authorisation.

Peer federations: Federations with which the EOSC AAI Federation has established transitive trust.

Registry: System used by the Federation Operator to manage entities and their metadata. This may be via a self-service tool or via other manual processes.

Registered representatives: Individuals authorised to act on behalf of the Member. These may take on different roles with different rights attached to them.

Relying Party: A service, site, or entity that depends on an identity provider to identify and authenticate a user who is requesting access to a digital resource.

[4] The Infrastructure proxy, as any other end service, can apply its own authorisation process concerning its protected resources. As such it may add Infrastructure specific attributes to the incoming identities. In other words, authorisation to access services may be based on a combination of information coming from more than one source, including the Infrastructure proxy.

3 EOSC AAI First Principles

Understanding and applying these principles are at the core of a successful, global, scalable and secure architecture for research and education AAI:

- User experience is the only touchstone.
- All trust flows from communities.
- There is no centre in a distributed system.

3.1 User experience is the first measure of success

Users, whether in the form of researchers, graduate students, teachers or administrators, are the reasons we build IT infrastructure of any kind. The authentication and authorization process is a necessary part of the security and risk management infrastructure of any IT service.

Security, while important, must always be deployed in conjunction with proper risk management to avoid over-engineering security controls at the expense of usability. We build technology for users, not for technologists. Since the process of authenticating to and obtaining the rights to use an IT service has no intrinsic value it must be made as unobtrusive as is humanly possible.

In EOSC we will adopt a scientific approach to usability and will quantify, measure and evaluate services from a usability perspective in order to be able to base our choices of technology, processes, user interface and service portfolios on empirical evidence about user behaviour. Quantitative user testing involves instrumenting services to allow data about user behaviour to be collected for later analysis. This approach is typically combined with incremental testing (aka A/B testing) to validate changes against stated goals for improvement to user journeys. In order to support a scientific approach to user testing, EOSC services and AAI in particular must be instrumented to support such a data-driven approach.

3.2 All trust flows from communities

Trust is the extent to which one party is willing to depend on the other party in a given situation with a feeling of relative security, even though negative consequences are possible[5].Trust does not derive from technology but is an emergent property of communities. A community in this context means a group of users, organised with a common purpose, and jointly granted access to resources. Communities may act as the interface between individual users and the resources. Trust in this sense means a community's ability to know and determine its users and their permissions. Technology can be used to represent certain aspects of trust (aka technical trust), typically by employing cryptography in various ways. Technical trust (the representation) is often confused with the thing itself. Trust is inherently local.

Trust is a property of certain well-established communities and as such cannot easily be made to scale beyond the limits of those communities. In EOSC we will build on the trust that exists within well-managed scientific collaborations and instead of trying to grow any one of those to a global context, we will create a framework within which many such communities can co-exist and interoperate.

[5] McKnight, H., Chervany, N. The Meanings of Trust. MISRC Working Paper Series 96-04, Management Information Systems Research Center, University of Minnesota, 1996

The goal of EOSC AAI is to provide the trust *mortar* with which we join the many bricks of the current set of scientific communities, collaborations and infrastructures together. For the user this means that what works today will work tomorrow, only better.

Interoperability between different community-based AAI solutions is achieved through the ongoing coordination work in the AEGIS group and the associated evolution of the Blueprint Architecture, as mentioned below.

While trust is defined as inherently local, establishing global trust across and between communities, research and e-infrastructure service providers is essential for the EOSC AAI to succeed. Defining this trust framework can be considered as one of the main challenges.

3.3 There is no centre in a distributed system

The term "EOSC AAI" has sometimes been interpreted as a singular instance of the EOSC AAI Architecture. Nothing could be further from the truth. The EOSC AAI is a set of principles and governance structures for how the architecture evolves and grows over time. As new science infrastructures emerge and evolve, each will find the best way to integrate into the EOSC ecosystem including the AAI.

Another way to state the same principle is this: The EOSC AAI is fully egalitarian in that no entity (organization, community etc) has power over any other. The functional EOSC AAI ecosystem will be based on the voluntary interoperable collaboration by all involved parties driven by enlightened self-interest.

Participants in the EOSC AAI will recognize that the objective is not to control users or to build walled gardens but to provide an open and fair playing field for service delivery to the scientific community.

3.4 A jump-off point for the EOSC AAI

The work to establish the architecture for the EOSC AAI needs a starting point. Fortunately, the AARC and AARC2[6] projects have provided just that. The AARC blueprint architecture[7] with all of its extensions and the ongoing governance within the AEGIS group provide just the starting point for the work ahead.

In particular the AARC BPA:

- builds on existing best practice in the scientific community
- provides clear guidance for how campus identity integrates with science
- comes with the beginnings of a governance structure in the AEGIS group[8]
- has international buy-in

Experience with the BPA in recent years has also identified an initial set of challenges that the EOSC AAI work will need to tackle, these are documented in the EOSC AAI Architecture Draft 2019.

[6] AARC and AARC2 projects: https://aarc-project.eu/

[7] AARC Blueprint Architecture: https://aarc-project.eu/architecture/

[8] AEGIS group: https://aarc-project.eu/about/aegis/

4 EOSC AAI Architecture

4.1 Introduction

The purpose of the EOSC AAI Architecture work is to establish a common global ecosystem for identity and access control infrastructures for the European Open Science Cloud (EOSC).

This chapter captures the current status of the EOSC AAI architecture discussions that are based on the AARC Blueprint Architecture 2019 and identify the challenges and the areas that require further work. The EOSC AAI Task Force has worked in collaboration with AEGIS, the AARC Community and other stakeholders to address these gaps, feed in the requirements from EOSC to the development of the AARC Blueprint Architecture 2020 and deliver the EOSC AAI Architecture 2020 version by the end of 2020.

As a starting point, this document is based on the AARC Blueprint Architecture 2019 (AARC-BPA-2019)[9]. The goal of the EOSC AAI is not to define a new AAI architecture, but rather to define an AAI architecture that follows the AARC BPA and the AARC Interoperability Guidelines and to work with the international community through AARC and AEGIS and shape the upcoming versions of the Blueprint Architecture to meet the evolving needs of EOSC.

4.2 The AARC Blueprint Architecture

The AARC Blueprint Architecture (BPA) provides a set of building blocks for software architects and technical decision makers who are designing and implementing access management solutions for international research collaborations.

AARC-BPA-2019 focuses on the interoperability aspects, to address an increasing number of use cases from research communities requiring access to federated resources offered by different research and e-Infrastructures. AARC-BPA-2019 introduces the concept of the Community AAI, which streamlines researchers' access to services. These typically include services offered to members of a specific community, as well as infrastructure services that may be shared with other communities. Users can authenticate to the Community AAI primarily via institutional credentials from national identity federations in eduGAIN, but, if permitted by the community, can also use other Identity Providers.

[9] https://doi.org/10.5281/zenodo.3672784

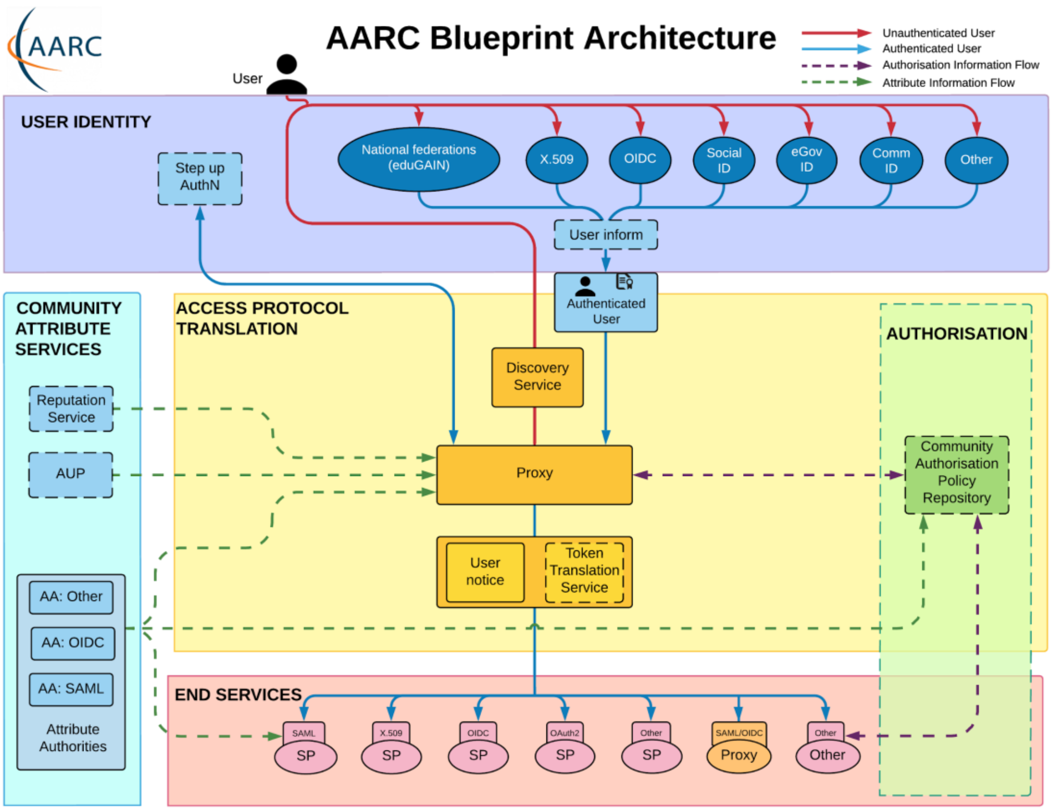

Figure 1: Component layers of the AARC Blueprint Architecture (AARC-BPA-2019)

AARC-BPA-2019 includes five (5) component layers: User Identity, Community Attribute Services, Access Protocol Translation, Authorisation, and End Services.

As illustrated in Figure 1, each of these layers includes one or more functional components, grouped by their complementary functional roles:

- *User Identity* - 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, more recently, 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.

- *Community Attribute Services* - Groups components related to managing and providing information (attributes) about users. Typically, this information includes community group memberships and roles, which is added on top of the information that might be provided directly by the identity providers from the User Identity Layer.

- *Access Protocol Translation*. 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 notice, which allows users to be informed regarding the processing of their personal data

- *Authorisation* - Contains components for controlling 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.
- *End-services* - 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.

4.3 AARC-BPA 2019 and the EOSC AAI

AARC-BPA 2019 distinguishes between two types of AAI services: One focuses on infrastructure management, while the other focuses on community management. Both types of AAI services may comprise the same interfaces (e.g., a proxy), but their functionality and their organisational purposes differ.

4.3.1 Community AAI

The purpose of the Community AAI is to streamline researchers’ access to services, both those provided by their own infrastructure (if they have one) as well as the services provided by infrastructures that are shared with other communities. 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. A critical part in the Community AAI is the Community Membership Management component. This functionality allows communities to register and manage their members in accordance with their own policies. Community membership and other roles and rights that are relevant for gaining access to services are recorded in the Community Attribute Services and exposed to services via a proxy or other forms of provisioning.

The Community AAI follows the proxy-based architecture shown in Figure 1. It can therefore add attributes to the federated identity that in turn can enable services to control access to their resources. Furthermore, the Community AAI is responsible for dealing with the complexity of using different identity providers with the services offered to the community. We can distinguish among three types of services (see also definitions in Chapter 2):

- *generic services* - provided to members of different communities, or individuals (e.g., the RCauth.eu Online CA service)
- *community services* - provided only to members of a given community
- *infrastructure services* - provided by a given research infrastructure or e-Infrastructure, typically through an infrastructure proxy

The architecture, from the perspective of a researcher is shown in Figure 2. This illustrates how community-specific services only need to connect to a single identity provider, i.e., their Community AAI. In contrast, generic services connect to multiple Community AAIs in order to serve different communities.

Being connected to multiple Community AAIs requires those generic services to provide some form of IdP discovery, to be able to redirect the user to the relevant Community AAI.

Additionally, to allow "community branding" of the service and automatically redirecting the user to the corresponding Community AAI, the generic services may support some means of doing "IdP hinting" (see [AARC-G049[10]]).

Communities may also require access to various services which themselves are behind (another) proxy. This could for example be resources offered by e-Infrastructures or Research Infrastructures (Infrastructures hereafter). These "Infrastructure Proxies" can be connected to multiple Community AAIs (see Figure 3). So, just as for the generic services, Infrastructure services should be able to hint to the Infrastructure Proxy which Community AAI to use (see [AARC-G049]).

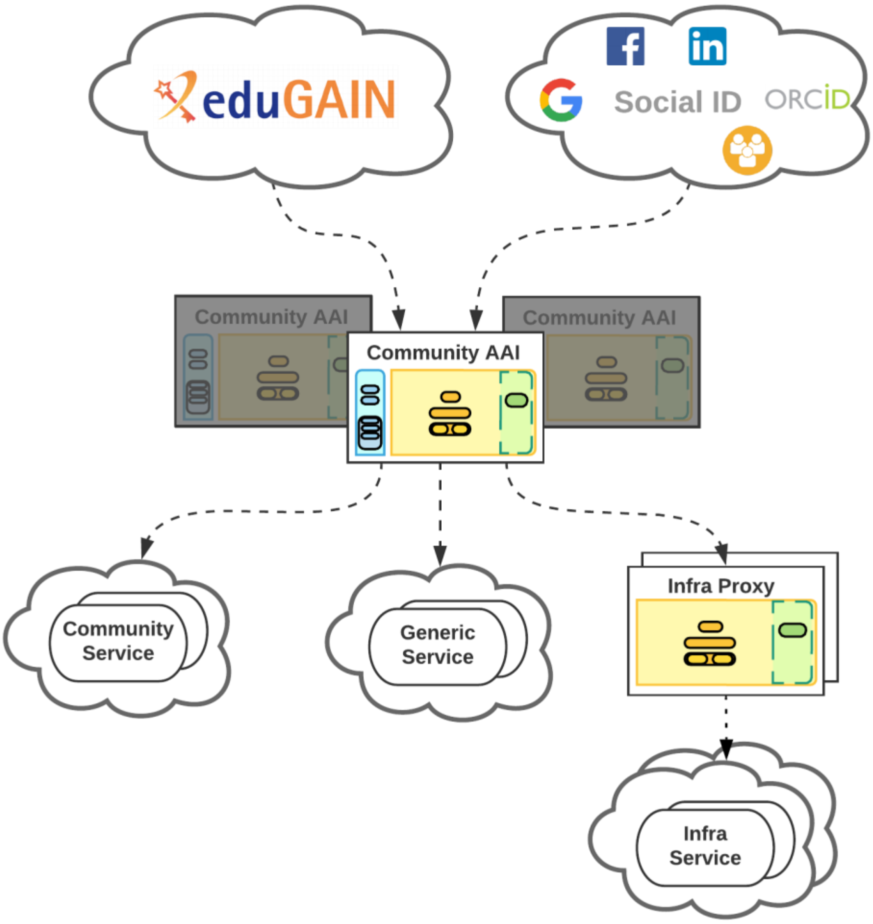

Figure 2: Researchers access resources through services using their institutional (eduGAIN), social or community-managed IdP via their Community AAI. Community services are connected to a single Community AAI, whereas generic services can be connected to more than one Community AAI. Infrastructure services are connected to different Community AAIs through a single infrastructure SP proxy.

It should be noted that this 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 is usually provided by a third-party infrastructure operator and can support multiple communities (tenants), as depicted in Figure 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

[10] https://aarc-project.eu/guidelines/aarc-g049/

scenario, the community managers are responsible for managing their community members, groups and authorisation attributes.

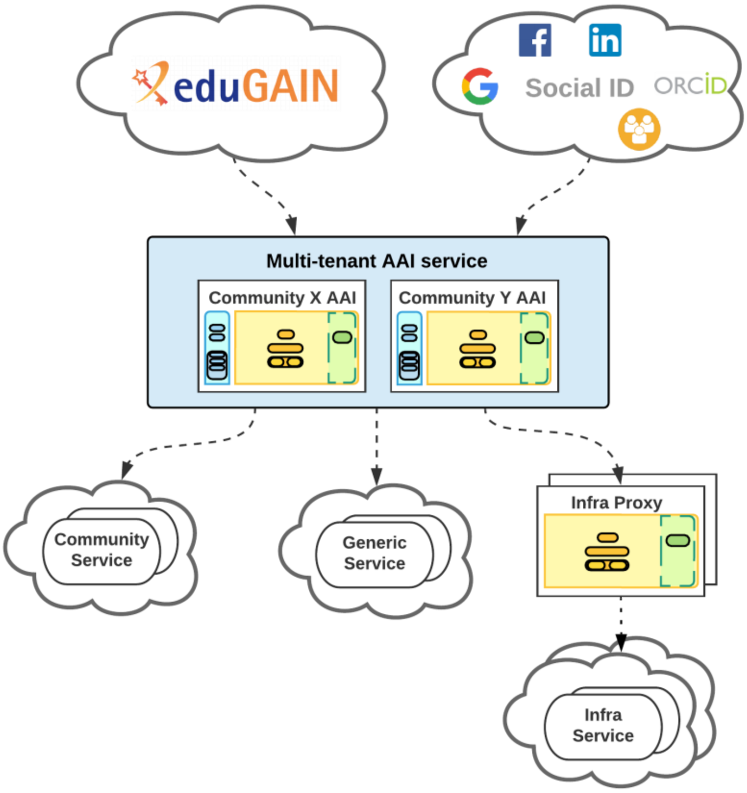

Figure 3: Multi-tenant deployment of AAI services following the AARC BPA architecture.

4.3.2 Infrastructure Proxy

The Infrastructure Proxy uses similar components and technologies as the Community AAI. However, the Infrastructure Proxy performs a very different role, as it is tasked to provide access to resources made available by an Infrastructure. The Infrastructure Proxy, enables Infrastructures with a large number of resources, to provide them through a single integration point, where the Infrastructure can maintain centrally all the relevant policies and business logic for making available these resources to multiple communities.

4.4 Challenges – Areas of further work

Below are a number of topics that were identified as needing addressing in AARC-BPA-2020 in order to meet the requirements of the EOSC AAI:

- The wording "community attribute services" is oversimplifying the important architectural component(s) of Attribute Providers (AtP) that can be independent from any IdP. It is true that in some cases AtPs and IdPs can be combined, but this should not be the general case that a blueprint must take into consideration. Beside the fact that "Community" is not well-enough defined as an authoritative entity, there is the general need for different - not only community-controlled - attribute/access management services. The general case should take aggregations of attributes originating from different AtPs into consideration; attributes that may not only be user-specific but also specific to other contexts and therefore are not necessarily managed by a "community" - Comment received from EUDAT.

- The current architecture works very well when the user is consuming services directly. One very common set of use cases though include the ability for a service agent to be able to act autonomously, on behalf of the user, consuming services and resources. If the services consumed by the agent are behind the same proxy, the current architecture works. For those cases, though, where an agent running on Service A needs to access resources on Service B, which might be connected by a different proxy, then there is no straight-forward solution at the moment. A solution for dynamically establishing trust in a distributed environment is provided by the OpenID Connect Federation implementer's draft[11]. The AARC community is working on AARC-G052: Recommendations for OpenID Connect/OAuth2 token-based access across different infrastructures[12] that is meant to be a temporary measure until OIDC-Fed is widely available.

- AARC-BPA-2019 introduced the logical separation between the Community AAIs and the Infrastructure Proxies with the assumption that the number of Community AAI can grow to a high number, but the number of the Infrastructure Proxies would remain low. This assumption is being challenged as in EOSC we are expecting that many Services will be made available from the national components of EOSC and probably, NRENs and National/European/Global Research Infrastructures and will be operating their own Infrastructure Proxies.

[11] http://openid.net/specs/openid-connect-federation-1_0.html

[12] https://docs.google.com/document/d/11Amv6kjPvVVgWB71iEaj6NcrhlNzht7HP9GJK6cNOS8/edit

5 EOSC AAI Federation Participation Policy Draft

5.1 Introduction and Applicability

This chapter describes the draft for a document documenting the structure and membership criteria of the *EOSC AAI Federation and that SHALL be published on the Federation website with a stable URL.* Such a document takes effect from the publication date shown on the cover sheet. All membership changes and entity registrations performed on or after that date SHALL be processed as described here until the document is superseded.

Requests to re-evaluate a given *member* or *entity* against a current version of this document MAY be made to the Federation Operator helpdesk.

This document SHOULD be considered in conjunction with the *EOSC AAI Federation Technical Interoperability Framework* and the EOSC AAI Policy and Security Baseline

5.2 Member Eligibility and Ownership

Members of the EOSC AAI Federation can be organisations providing Services to the EOSC; organisations operating Proxies that aggregate other service providers or enrich identities; and providers of authentication Identities whose identities are used by EOSC services or by EOSC proxies. Membership to the EOSC AAI Federation shall be open to organisations that have a relationship with the EOSC, wholly or in part.

Members MUST abide by the EOSC Rules of Participation, and maintain compliance with these Rules for as long as they have entities registered in the EOSC AAI Federation Registry. In addition, they MUST meet and maintain compliance with the EOSC AAI Federation organisational baselines specified hereunder.

Members can register entities in the EOSC AAI Federation Registry and have their metadata published in the EOSC AAI Federation metadata. Members MUST enter into an understanding with the EOSC AAI Federation Operator, and SHOULD execute the understanding through a legal entity. If the member is a (research) community, the requirement for a legal entity MAY be waived, under the condition that (i) the community has established between its members a framework of agreements, or a memorandum of understanding, between legal entities that describes the criteria for participation in the community, and (ii) the community is established on a long-term basis, and (iii) the community designates two individuals who are authorized to represent the community towards the EOSC AAI Federation Operator.

Members MAY also be members of a peer federation of the EOSC AAI Federation. In such a case, they SHALL NOT register any entities that are already registered in a Peer Federation.

Membership of the EOSC AAI Federation MUST be requested to the Federation Operator by each prospective member. In this request, the applicant MUST:

- declare its intent to join the EOSC AAI Federation;
- declare its participation in the EOSC and adherence to its Rules of Participation;
- commit to adherence to the pertinent technical requirements of the EOSC AAI Interoperability Framework (technical baseline);
- commit to adherence to the security policy baseline of EOSC security operations;
- provide contact information for administrative, technical, and security matters, each of which *Registered Representatives* SHALL have least two contact entry points;

- provide a preferred name and identifier for the organisation, in the understanding that this name serves as a non-binding suggestion to the EOSC AAI Federation Operator for any name assignments;

- commit to keeping the provided information accurate and up-to-date.

The EOSC AAI Federation Operator SHALL verify information provided by the applicant using verifiable data sources, including data sources maintained by the EOSC Organisation.

Importing entities from Peer Federations is based on the transitive trust between those Federations and the EOSC AAI Federation. Imported entities need to meet the pertinent technical requirements of the EOSC AAI Interoperability Framework (technical baseline) and with the security policy baseline, and for which the requisite contact information can be obtained. Either the organisation directly, or the peer federation from which its entities are imported SHALL commit to keeping the information accurate and up-to-date. Imported entities SHALL NOT be re-exported to any peer federations.

Entities acting as Relying Parties SHALL NOT be imported unless the organisation responsible for the entity has declared their intent to participate in the EOSC with these relying parties, and to adhere to the EOSC Rules of Participation with respect to these relying parties.

The EOSC AAI Federation Operator SHOULD remove an organisation and all entities of that organisation from the federation, once that organisation no longer meets the criteria for membership. Exceptionally, the EOSC AAI Federation Operator MAY grant a grace period for removal from the federation at its own sole discretion, if such an exception serves the overarching goal of a trustworthy and functional EOSC. Such an exception MUST be time-limited, documented, and approved by the highest level of management commensurate with the impact of the continued federated availability of the entities involved.

In the pursuit of EOSC accessibility and coverage, the EOSC AAI Federation Operator MAY import into the federation from any source, entities that provide Identity authentication, or that act as a Proxy source of identity information, subject to conditions of trustworthiness and technical compliance. Such entities MUST be imported from a trusted source, and MUST meet the pertinent requirements of the EOSC AAI Federation Technical Framework. The EOSC AAI Federation Operator MAY indicate such entities through specific metadata attributes or specific policy URIs, depending on its technical capabilities. Organisations of which such entities are discretionarily imported SHALL NOT thereby become Members of the EOSC AAI Federation, nor have the obligation to meet its requirements. At its own discretion, the EOSC AAI Federation Operator MAY remove such entities at any time, without warning and without redress.

Members MAY withdraw from the EOSC AAI Federation by informing the Operator. The Operator MUST verify the identity of the requester(s) and their authority to request termination of membership. Termination of membership of the EOSC AAI Federation does not release organisations from any prior commitments that are intended to survive termination.

6 EOSC AAI Federation Technical Framework Draft

6.1 Introduction and Applicability

This draft document describes the entity registration practices of the EOSC AAI Federation Operator with effect from the publication date shown on the cover sheet. For an entity to be eligible to be registered in the federation, its organisation must meet the eligibility criteria of the federation as defined in the EOSC AAI Federation Participation Policy as described in the previous chapter and the entity must adhere to the technical baseline as laid out in this chapter. All new entity registrations performed on or after that date SHALL be processed as described here until the document is superseded.

The document SHALL be published on the Federation website at a stable URL.

Metadata about an entity that does not include a reference to a registration policy MUST be assumed to have been registered under an historic, undocumented registration practice regime. Requests to re-evaluate a given entity against a current version of this document MAY be made to the EOSC AAI Federation Operator helpdesk/registration tool

6.2 Entity operational and technical Requirements

6.2.1 Entity Requirements

The EOSC AAI Federation Operator SHALL import entities of service providers only from those organisations that have indicated their participation in the EOSC AAI by applying for Membership as described in the EOSC AAI Federation Participation Policy. It SHALL NOT import entities from organisations that do not yet - or do no longer - meet the eligibility requirements, nor will it import entities that do not meet the validation requirements defined in this document.

In the pursuit of EOSC accessibility and coverage, the EOSC AAI Federation Operator MAY import into the federation from any source entities that provide Identity authentication, or that act as a Proxy source of identity information, subject to conditions of trustworthiness and technical compliance. Such entities MUST be imported from a trusted source, and MUST meet the pertinent requirements defined in this document. The EOSC AAI Federation Operator MAY indicate such entities through specific metadata attributes or specific policy URIs, depending on its technical capabilities. At its own discretion, the EOSC AAI Federation Operator MAY remove such entities at any time, without warning and without redress.

Entities in the EOSC AAI Federation MUST implement and comply with the AARC Interoperability Guidelines approved by AEGIS [AARC-Guidelines[13]].

6.2.2 Entity Registration

The process by which a Federation member can register an entity is described at <url-to-be-defined>.

The Federation Operator SHALL require the following information for each entity registration application and for any entity imported from Peer Federations:

- Entity name (in English and optionally in other languages supported by the entity)
- Entity description

[13] https://wiki.geant.org/display/AARC/AARC+Interoperability+Guidelines+Approved+by+AEGIS

- Website URL for information about the entity; the content found at the URL SHOULD provide more complete information than what is provided in the description
- Member information including:
    - Name of the member organisation
    - Display name of the member organisation used for user facing interfaces
    - Website URL for information about the organisation
- Member contact information of the following types:
    - Technical and/or Helpdesk/Support contact information (for redirecting users)
    - Security/incident response (see also Sirtfi)
    - Administrative (optional)
- Entity contact information of the following types: (if different from the member's contact information)
    - Technical and/or Helpdesk/Support contact information (for redirecting users)
    - Security/incident response (see also Sirtfi)
    - Administrative (optional)
- Whether a Logo URL is provided (for showing in catalogues); if provided, logos SHOULD:
    - use a transparent background where appropriate to facilitate the usage of logos within a user interface
    - use PNG, or GIF (less preferred), images
    - use HTTPS URLs in order to avoid mixed-content warnings within browsers
    - have a size smaller than 50000 characters when encoded in base64
- For entities acting as Identity providers:
    - The entity SHOULD support the REFEDS Research & Scholarship Entity Category [REFEDS-R&S[14]]
    - The entity MUST release at least the following attributes (please refer to Annex II for more details):
        - A non-reassignable, persistent, unique user identifier
        - Name information
        - Email information

[14] https://refeds.org/category/research-and-scholarship

        - in the case of an academic identity provider, the entity MUST release
            - Home Organization information
            - Affiliation information
    - The entity is compliant with the Security Incident Response Trust Framework for Federated Identity (Sirtfi) [REFEDS-Sirtfi[15]]

- For entities acting as Relying Parties:
    - Link to the privacy policy
    - Link to the Acceptable Use Policy / Terms of Use
    - The entity is compliant with the GEANT Code of Conduct version 1 [GEANT-DPCoCo[16]] or any other code of conduct compatible with legislation and guidelines on data protection and privacy including GDPR
    - Whether the entity is compliant with the REFEDS Research & Scholarship Entity Category [REFEDS-R&S]
    - The entity is compliant with the Security Incident Response Trust Framework for Federated Identity (Sirtfi) [REFEDS-Sirtfi]
- The federated protocol(s) the entity is supporting

Other technical information related to the supported federated protocol(s)

6.3 SAML entity technical requirements

Identity Providers and Service Providers supporting the SAML protocol MUST comply with the SAML2int SAML 2.0 Interoperability Deployment Profile [SAML2int[17]].

6.3.1 Metadata

Identity Providers and Service Providers supporting the SAML protocol MUST provide a SAML 2.0 Metadata document representing its entity according to the SAML2int SAML 2.0 Interoperability Deployment Profile [SAML2int] and the recommendations for upstream metadata produced by eduGAIN participants [eduGAIN-Metadata-Recommendations[18]].

6.3.2 Entity Identifier Format

For entities supporting the SAML protocol, the entity identifier is the SAML entityID attribute. Values of the entityID attribute registered MUST be an absolute URI using the http, https or urn schemes.

https-scheme URIs are RECOMMENDED for all entities.

http-scheme and https-scheme URIs used for entityID values MUST contain a host part whose value is a DNS domain.

[15] https://refeds.org/sirtfi

[16] https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home

[17] https://kantarainitiative.github.io/SAMLprofiles/saml2int.html

[18] https://wiki.geant.org/display/eduGAIN/Best+Current+Practice

6.3.3 Scope Format

For Identity Provider entities supporting the SAML protocol, scopes MUST be rooted in the DNS domain namespace, expressed in lowercase. Multiple scopes are allowed.

Regular expressions representing multiple scopes MAY be used, but all DNS domains covered by the expression SHALL be included in checks by the Federation Operator for the member's right to use those domains. For these checks to be achievable by the Federation Operator, the set of DNS domains covered by the regular expression MUST end with a domain under a public suffix – that is, a literal '.', followed by at least two DNS labels separated by literal '.' (representing a domain to be validated as "owned" by the entity owner), and ending with a '$' anchor (e.g. (foo | bar)\.example\.com$).

6.4 OpenID Connect / OAuth entity technical requirements

The requirements defined in this chapter apply to entities implementing the OpenID Connect protocol. To support user access through SAML-based Identity Providers, it is assumed that OpenID Connect relying parties will be connected to AARC-compliant [AARC-Guidelines[19]] IdP-SP-Proxy services registered in the EOSC AAI Federation that are able to translate between SAML and OpenID Connect.

6.4.1 Metadata

Identity Providers supporting OpenID Connect MUST describe their configuration through a well-known location as a JSON document following the [OIDC-Discovery[20]] specification.

Relying parties supporting OpenID Connect SHOULD retrieve the Identity Provider's configuration based on the Issuer information using the [OIDC-Discovery] specification.

6.4.2 Grant types

OpenID Connect relying parties utilising the authorization grant type SHOULD use PKCE [RFC7636[21]] in conjunction with the authorisation server in order to detect and prevent attempts to inject (replay) authorisation codes into the authorisation response. The PKCE challenges must be transaction-specific and securely bound to the user agent in which the transaction was started. OpenID Connect relying parties MAY use the "nonce" parameter of the OpenID Connect authentication request as specified in [OIDC-Core] in conjunction with the corresponding ID Token claim for the same purpose. It is NOT RECOMMENDED to use the implicit grant and any other response type causing the authorization server to issue an access token in the authorization response.

6.4.3 Claims

OpenID Connect relying parties MUST support requesting Claims about the End-User and the Authentication event using specific scope values as described in [OIDC-Core[22]]. Requesting individual Claims using the claims request parameter as defined in [OIDC-Core] SHOULD also be supported.

Claims which are not part of the standard set of claims defined in [OIDC-Core] SHOULD be requested following the mapping recommendations described in [OIDC-REFEDS].

[19] https://wiki.geant.org/display/AARC/AARC+Interoperability+Guidelines+Approved+by+AEGIS

[20] https://openid.net/specs/openid-connect-discovery-1_0.html

[21] https://tools.ietf.org/html/rfc7636

[22] https://openid.net/specs/openid-connect-core-1_0.html

6.4.4 Redirection URIs

OpenID Connect relying parties MUST pre-register one or more Redirection URI to which authentication responses from the OpenID Connect Identity Provider will be sent. The Identity Provider MUST utilise exact matching of the redirect URI specified in an authentication request against the pre-registered URIs [OAuth2-BCP[23]], with the matching performed as described in [RFC3986[24]] (Simple String Comparison). Redirection URIs MUST use the schemata defined in Section 3.1.2.1 of the [OIDC-Core] specification.

6.5 Entity Management

The Registered Representatives for a member of the EOSC AAI Federation MAY add, modify or remove any number of entities that meet the requirements specified in this document.

6.5.1 Entity Change Requests

Any request for entity addition, change or removal from members of the EOSC AAI Federation needs to be communicated from or confirmed by the Registered Representatives for the entity.

For the purpose of communicating changes, the EOSC AAI Federation may provide/mandate the use of a Registry tool.

6.5.2 Unsolicited Entity Changes

The EOSC AAI Federation Operator MAY amend, modify or remove the metadata of an entity at any time in order to:

- Ensure the security and integrity of the EOSC AAI Federation
- Ensure the security and integrity of the metadata;
- Comply with agreements with Peer Federations;
- Improve interoperability;
- Add value to the metadata.

Changes will be communicated to Registered Representatives for the entity without undue delay.

6.5.3 EOSC AAI Federation Operator initiated changes requests

In some cases, the EOSC AAI Federation Operator MAY request members to make updates to their entities metadata in order to:

- Improve compliance with the EOSC AAI Federation technical baseline;
- Ensure the security and integrity of the metadata;
- Comply with agreements with Peer Federations;
- Improve interoperability;

[23] https://tools.ietf.org/id/draft-ietf-oauth-security-topics-16.txt

[24] https://tools.ietf.org/html/rfc3986

In such cases, the EOSC AAI Federation Operator MUST provide a justification for the requested change and a time limit by when members MUST have implemented the change. This limit SHALL be at most six months from the publication of the change request, but may be less if justified by e.g., security considerations.

If the Member fails to implement the change request in time, the relevant entity(ies) MAY be temporarily or permanently removed from the federation.

The EOSC AAI Federation Operator MUST remove all entities of an organisation from the federation, if that organisation no longer meets the criteria for participation defined in EOSC AAI Federation Participation Policy or that organisation departs from the EOSC AAI Federation. The EOSC AAI Federation Operator MUST remove entities that no longer meet the criteria for participation defined in the EOSC AAI Federation Participation Policy or those entities no longer comply with the technical and operational requirements listed in this document. The EOSC AAI Federation Operator MAY grant a grace period for removal of an entity from the federation at its own sole discretion, if such an exception serves the overarching goal of a trustworthy and functional EOSC. Such an exception MUST be time-limited, documented, and approved by the highest level of management commensurate with the impact of the continued federated availability of the entities involved.

6.6 Entity Validation

For each new entity registration application, the Federation Operator SHALL carry out entity validations checks. These checks include:

- Ensuring the member’s right to use particular domain names in relation to protocol endpoints, user facing URL endpoints, entityID attributes and, for Identity Provider entities, any scope elements.

  The right to use a domain name SHALL be established in one of the following ways:

    - By means of the approved domain validation methods for public web trust
    - A member MAY be granted the right to make use of a specific domain name through a permission letter from the domain owner on a per-entity basis. Permission SHALL NOT be regarded as including permission for the use of sub-domains.

- Ensuring the member’s right to use URN namespaces;
- Ensuring all required information is present in the metadata provided or the entity registration application;
- Ensuring protocol endpoints provide correct information;
- Ensuring protocol endpoints are properly protected with TLS / SSL certificates;
- Ensure user facing URL endpoints are available
- Ensure technical, administrative and security contact information is periodically validated

7 Annex I: Metadata Format

Metadata for all entities registered by the Federation Operator SHALL make use of the [SAML-Metadata-RPI-V1.0] metadata extension to indicate that the Federation Operator is the registrar for the entity and to detail the version of the MRPS statement that applies to the entity. The following is a non-normative example:

```
<mdrpi:RegistrationInfo

  registrationAuthority=”http://federation.org”

  registrationInstant=”2020-11-29T13:39:41Z”>

  <mdrpi:RegistrationPolicy xml:lang=”en”>

      https://federation.org/doc/MRPS20200924

  </mdrpi:RegistrationPolicy>

</mdpri:RegistrationInfo>
```

8 ANNEX II: ATTRIBUTE FORMAT

User identity attribute types expressed as SAML attributes or OpenID Connect claims:

Identity Attribute Type	SAML Attribute	OpenID Connect Claim
Non-reassignable, persistent, unique user identifier	- subject-id [SAML-SubjectID-v1.0] - pairwise-id [SAML-SubjectID-v1.0] - SAML Persistent NameID [] - eduPersonPrincipalName [eduPerson-202001] - eduPersonTargetedID [eduPerson-202001] - eduPersonUniqueId [eduPerson-202001] - voPersonID [voPerson-2.0]	- sub (public) [OIDC-Core] - sub (pairwise) [OIDC-Core] - eduperson_targeted_id [eduPerson-202001] - eduperson_unique_id [eduPerson-202001] - voperson_id [voPerson-2.0]
Name information	Either (or all) of the following: - cn - displayName - givenName + sn	Either (or all) of the following: - name [OIDC-Core] - given_name + family_name [OIDC-Core]
Email information	Either (or both) of the following: - mail - voPersonVerifiedEmail [voPerson-2.0]	Either (or all) of the following: - email [OIDC-Core] - email_verified [OIDC-Core] - voperson_verified_email [voPerson-2.0]
Home organisation information	schacHomeOrganization [SCHAC-1.5]	schac_home_organization [SCHAC-1.5]
Affiliation	Either (or both) of the following: - eduPersonScopedAffiliation [eduPerson-202001] - voPersonExternalAffiliation [voPerson-2.0]	Either (or both) of the following: - eduperson_scoped_affiliation [eduPerson-202001] - voperson_external_affiliation [voPerson-2.0]

The EOSC Architecture Working Group has assigned the Authentication and Authorization Infrastructure Task Force (AAI TF) the task to establish a common global ecosystem for identity and access control infrastructures for the European Open Science Cloud (EOSC).

This document is the final deliverable of the EOSC AAI Task Force, operating in the context of the EOSC Architecture Working Group and will be handed over to the EOSC Association.

Research and Innovation policy

Publications Office
of the European Union