EOSC Future

D7.3

EOSC Federated Authorisation and Authentication Activities

The *EOSC Future* project is co-funded by the European Union Horizon Programme call INFRAEOSC-03-2020, Grant Agreement number 101017536

Version 1.0
February 2023

D7.3 / EOSC Federated Authorisation and Authentication Activities

Lead by GÉANT
Authored by Christos Kanellopoulos (GÉANT), Nicolas Liampotis (GRNET)
Reviewed by Klaas Wierenga (GÉANT), Mark Dietrich (EGI) and Athanasia Spiliotopoulou (JNP)

Dissemination Level of the Document

Public

Abstract

This report presents the activities taking place in the AAI task (7.3) from the beginning of the project until M20. The document goes over the work in international standardisation activities, the EOSC-Core Infrastructure Proxy, RCAuth CA and the IGTF Certificate Proxy, the AAI Fabric Monitoring, the EOSC AAI Federation and the coordination with the Science Cluster, Research Infrastructures, and e-Infrastructures.

Version History

Version	Date	Authors	Description
V0.1	10/10/2022	Christos Kanellopoulos (GÉANT)	Initiation
V0.2	18/11/2022	Christos Kanellopoulos (GÉANT)	TOC
V0.3	14/12/2022	Nicolas Liampotis (GRNET)	RCAuth CA
V0.4	21/12/2022	Nicolas Liampotis (GRNET)	Standardisation activities
V0.5	30/12/2022	Nicolas Liampotis (GRNET)	Support Services
V0.6	13/01/2023	Christos Kanellopoulos (GÉANT)	EOSC AAI Federation
V0.7	17/01/2023	Christos Kanellopoulos (GÉANT)	Science Clusters, RIs, eInfrastructures
V0.8	18/01/2023	Christos Kanellopoulos (GÉANT)	All sections completed
V0.9	25/01/2023	Christos Kanellopoulos (GÉANT)	Review Comments incorporated
V1.0	03/02/2023	Christos Kanellopoulos (GÉANT)	Final Version to be submitted to EC by PC (ATHENA)

Copyright Notice

Table of Contents

Table of Tables

Table of Figures

Glossary

EOSC Future project Glossary is incorporated by reference: https://wiki.eoscfuture.eu/x/JQCK

List of Abbreviations

Acronym	Definition
AAI	Authentication & Authorization Infrastructure
AARC	Authentication and Authorization for Research and Collaboration
BPA	Blueprint Architecture
CA	Certification Authority
CUID	Community User Identifier
DS	Delegation Server (component of RCauth Online CA)
OIDC	OpenID Connect
SAML	Security Assertion Markup Language
WAYF	Where Are You From (also known as Identity Provider Discovery Service)

1 Executive Summary

In 20 months since the beginning of the project, the AAI task was able to deliver on the promise of the EOSC Federated AAI. The focus of the task has been the implementation of the EOSC AAI following the recommendation of the EOSC Authentication and Authorisation Infrastructure report published by the European Commission, Directorate-General for Research and Innovation in January 2021 [EOSC-AAI-2021].

The work took place in the following areas:

- Improving the EOSC-Core Infrastructure Proxy and providing operational support
- Adding support for High Availability in the RCAuth CA
- Ensuring the operation of the IGTF x509 to SAML2 Bridge
- Implementing support tools to ensure the quality of the delivered services
- Implementing the EOSC AAI Federation
- Supporting the Science Clusters, Research Infrastructures and e-Infrastructures to use the EOSC Federated AAI
- Contribute to international standardisation activities.

In addition, the task has been actively contributing to the design of the EOSC AAI Architecture v2022 report, which will further guide the implementation of the EOSC Federation in 2023 onwards.

2 Introduction

This report presents the work performed by the AAI task in Work Package 7 of the EOSC Future Project starting from the beginning of the project until October 2022.

The goal of the AAI task is to plan, deliver and continuously improve the EOSC Federated AAI. The EOSC Federated AAI will enable Service Providers to deliver services and access to resources to research communities and individual researchers, allowing users to use their institutional, community and eIDAS enabled digital identities.

The task is responsible for:

- Planning, delivering and continuously improving the EOSC Federated AAI under definition within the WG ARCH of EOSC. This includes the AARC compliant proxies for EOSC-Core Services as well as the auxiliary services for the EOSC AAI and a white labelled online certification authority (RCauth).

- Integrating Community AAIs, Identity Providers and Services into the EOSC Federated AAI and providing coordinated consultancy, support and guidelines on how to connect Service Providers, Identity Providers and Community AAI Services to the EOSC Federated AAI.

- Managing incidents and service requests for the AAI services, restoring normal/agreed service operation within the agreed time after the occurrence of an incident, and responding to user service requests, namely user requests for information, advice, access to a service or a pre-approved change.

- Implementing change control for the AAI services, to ensure changes to configuration items are planned, approved, implemented and reviewed in a controlled manner to avoid adverse impact of changes to services or the customers receiving services. It is noted that task participants will be participating and contributing to the work in international fora, such as AEGIS, AARC Community and FIM4R.

3 International Standardisation and Harmonisation activities

During the reporting period, the task participated in AEGIS, the AARC-Community Working Groups and the AAI Task Force of the EOSC Association, contributing to the following standardisation and harmonisation activities:

- **Guidelines for expressing community user identifiers (AARC-G026)**: Defines how to express Community User Identifiers (CUIDs) such that the identifier values can be transported in an interoperable way across AARC BPA-compliant AAI services. The CUID is a subject identifier, where the subjects are generally but not exclusively natural persons. The CUID is an attribute of the subject's digital identity which is managed by the Community AAI. The guidelines specify how the CUID is communicated from the Community AAI to its connected services, i.e. infrastructure services (accessible through Infrastructure Proxies), generic and community services.

- **A specification for hinting an IdP which discovery service to use (AARC-G062)**: Defines the `aarc_ds_hint` hint which can be used for routing the user to a specific Discovery Service for selecting their preferred Identity Provider.

- **A specification for providing information about an end service to a Discovery Service (AARC-G063)**: Defines the `aarc_service_hint` hint which can be passed onto the Discovery Service to present the user with information about the service that requested authentication.

- **Guidelines for expressing group membership and role information (AARC-G069)**: Defines a URN namespace for expressing group membership and role information across AARC BPA-compliant AAI services, using common identity federation protocols, namely SAML and OpenID Connect/OAuth 2.0.

- **AARC Community-based Access Entity Category (AARC-G079)**: Defines an entity category to support the release of attributes to Service Providers that have a proven need to receive a set of community-managed information about their users in order to effectively provide their service to the users. A final draft of the document is under public consultation.

- **OAuth 2.0 Proxied Token Introspection (AARC-G052)**: Extends the OAuth 2.0 Token Introspection (RFC7662) method to allow conveying meta-information about a token from an OAuth 2.0 Authorization Server to the protected resource even when there is no direct trust relationship between the protected resource and the token issuer. A final draft of the document is under public consultation.

- **EOSC AAI Architecture v2022**: This version of the EOSC AAI Architecture builds on the work performed by the AAI task force of the Architecture Working that was published in January 2021. The 2022 version of the EOSC AAI Architecture focuses on the following elements:
  - consistent user experience and interfaces for service providers
  - multi-infrastructure workflows scaling trust
  - growth of EOSC beyond the research and education community
  - user and community attributes and authorisation.

  The EOSC AAI Architecture v2022 is expected to be published by the EOSC Association in January 2023.

4 EOSC-Core Infrastructure Proxy

The EOSC-Core Infrastructure Proxy is an AAI service that connects the EOSC-Core and EOSC Support Services to the EOSC Federated AAI. Services that are eligible for connecting to the EOSC-Core Infrastructure Proxy are:

- EOSC-Core Services
- EOSC Support Services operated in the context of the EOSC Future project
- EOSC Support Services operated by the EOSC Secretariat and/or the EOSC Association

The EOSC-Core Infrastructure Proxy implements the AARC Blueprint Architecture and interoperability guidelines. Specifically, it acts as an Infrastructure Proxy allowing users to use their institutional, social, ORCID, and community digital identities for accessing EOSC-Core and EOSC Support services.

All critical service components of the EOSC-Core Infrastructure Proxy are operated in High Availability mode. The deployment architecture can scale horizontally by provisioning more nodes, if required to increase service capacity. During the reporting period, the backend database store was configured to operate in clustered mode, supporting streaming replication and Point-in-Time Recovery for a period of six months in the past. Apart from the production instance, there is a demo/acceptance instance (also operated in High Availability mode) that allows for testing the integration with new services without affecting the production environment.

Since the beginning of the project, the following services were connected to the EOSC-Core Infrastructure Proxy in production:

- EOSC Collaborations tools:
  - EOSC WIKI (Confluence)
  - EOSC Issue Tracker (Jira)
- EOSC Helpdesk
- EOSC Open Science Observatory
- EOSC Topology (GOCDB)
- EOSC Search Service.

Furthermore, the demo instance of the EOSC Accounting for Services was connected to the demo instance of the EOSC-Core Infrastructure Proxy.

During the reporting period, the EOSC-Core Infrastructure Proxy implemented the following AARC interoperability guidelines:

- Guidelines for expressing community user identifiers (AARC-G026) using the voPersonID attribute in SAML or the voperson_id Claim in OpenID Connect.
- Guidelines for expressing affiliation information (AARC-G025) using the voPersonExternalAffiliation attribute in SAML or the voperson_external_affiliation Claim in OpenID Connect.

Furthermore, the operations team addressed issues related to the release of insufficient user attributes from a number of Identity Providers in eduGAIN. Resolving such issues required communication with the administrators of the Identity Providers in order to configure their attribute release policy for the EOSC-Core Infrastructure Proxy to include the attributes defined in the REFEDS Research and Scholarship attribute bundle. All but one of the reported user attribute release issues were resolved, as detailed in Table 4.1.

Table 4-1: Resolution status of issues related to the release of user attributes from Identity Providers

Identity Provider	Resolution	Support channel
Max Planck Institute for Plasma Physics (IPP)	Resolved on 10 Oct 2022	EOSC Helpdesk Ticket #1392
BBMRI	Resolved on 27 Sept 2022	EOSC Helpdesk Ticket #1446
University of Manchester	Resolved on 27 July 2022	mail
University of Limerick	Resolved on 7 July 2022	EOSC Helpdesk Ticket #673
Tampere Universities	Resolved on 4 July 2022	EOSC Helpdesk Ticket #783
University of Reading	Resolved on 8 June 2022	mail
Helmholtz-Zentrum Dresden-Rossendorf e.V. (HZDR)	Resolved on 26 Apr 2022	mail

University College Dublin	Resolved on 17 Mar 2022	Legacy EOSC Helpdesk Ticket #1951
Helmholtz-Zentrum Berlin für Materialien und Energie GmbH (HZB)	Resolved on 14 Feb 2022	Legacy EOSC Helpdesk Ticket #1939
Bangor University	Resolved on 17 Dec 2021	Legacy EOSC-Core Infrastructure Proxy Helpdesk Ticket #155163
Reykjavík University	Closed as unresolved on 17 Dec 2021 due to lack of response from the Identity Provider support	Legacy EOSC-Core Infrastructure Proxy Helpdesk Ticket #155017
ESO - European Southern Observatory	Resolved on 2 July 2021	mail
King's College London	Resolved on 8 June 2021	mail
Maynooth University	Resolved on 13 July 2020	mail

It should be noted that the operation of the EOSC-Core Infrastructure involved technological upgrades of the underlying OpenID Provider framework to improve compliance with the OpenID Connect specification and OAuth 2.0 Best Current Practices, as well as to add support for the OAuth 2.0 Proxied Token Introspection. The upgrade has been completed in the demo/acceptance environment, while deployment in production is pending.

The number of monthly user logins through the EOSC-Core Infrastructure Proxy is illustrated in Figure 4.1.

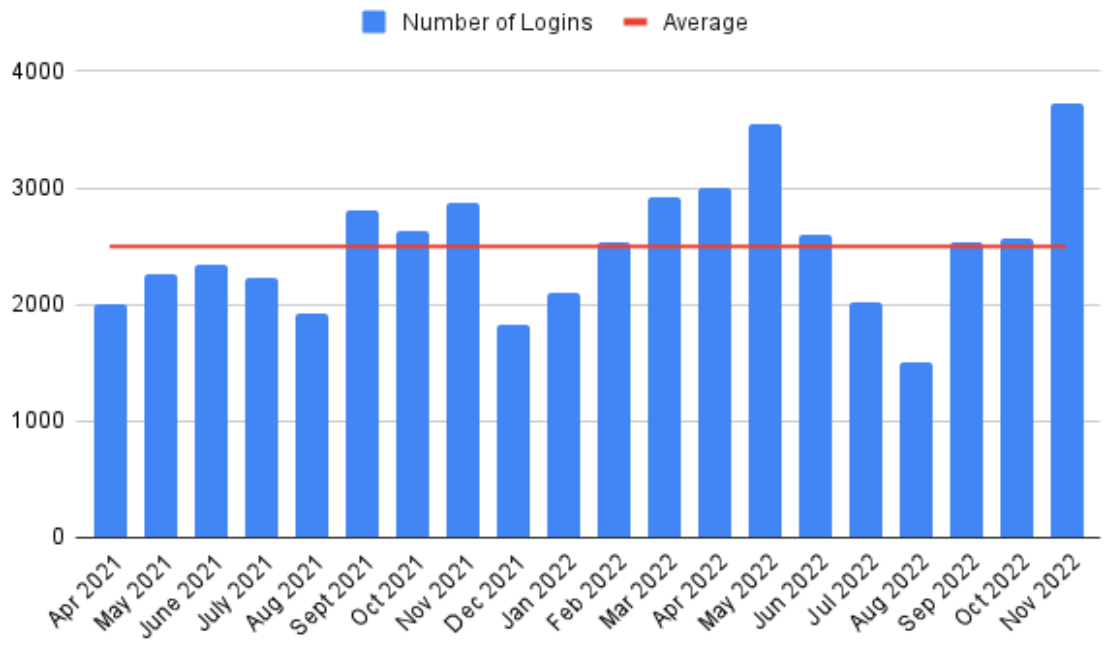

Figure 4.1: Monthly number of user logins through the EOSC-Core Infrastructure Proxy

The cumulative sum of unique users accessing services through the EOSC-Core Infrastructure Proxy is illustrated in Figure 4.2.

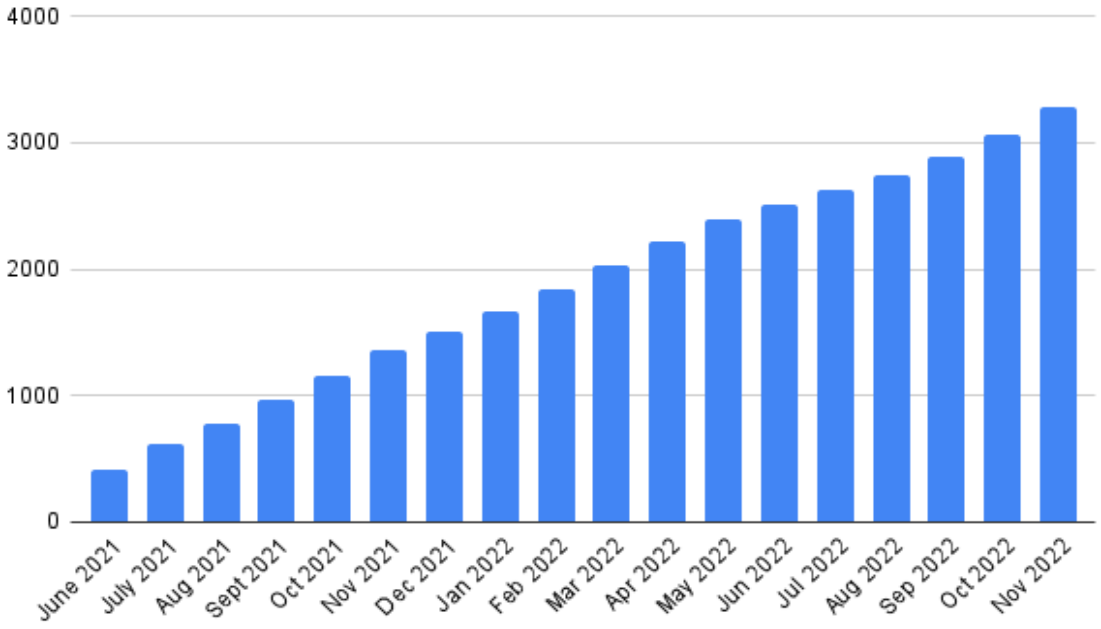

Figure 4.2: Cumulative sum of unique EOSC-Core Infrastructure Proxy users per month

5 RCauth Certification Authority

The focus of this activity in the AAI task was to introduce High Availability capabilities in the infrastructure supporting the RCAuth CA.

RCauth comprises three main service components: a database, a signing server, and a web server (delegation server, DS) through which the clients interact with the CA. It also uses a filtering Identity Provider Discovery (Where Are You From - WAYF) service to connect to eduGAIN, accepting only Identity Providers that assert both the REFEDS Research and Scholarship Entity Category and the Sirtfi security framework. In addition, several other Identity Providers, including Community AAIs (e.g. Check-in and eduTEAMs), are accepted following an individual policy agreement and exchange of metadata.

Originally just running as a single instance at Nikhef, RCauth CA is now operated in three distinct physical locations: at Nikhef in Amsterdam, The Netherlands; at STFC's RAL site in South Oxfordshire, UK; and at GRNET in Athens, Greece. Each of these operating locations has specific facility, management, and operational controls that are described separately in the RCauth ICA CP/CPS.

During the reporting period, both RAL and GRNET deployed production instances of the RCauth service components fully replicating the production instance at Nikhef. To this end, a distributed database was set up to retain near-synchronous state across all instances over transport-protected L4 virtual circuits.

Furthermore, the private key used in production was replicated from Nikhef to RAL and GRNET following a secure exchange process using one-time pads (OTP), a key exchange which offers perfect security if executed correctly. OTP requires splitting the key into parts and secure exchange of these parts individually. As most of the exchange of parts had been done before the pandemic by exchanging secrets at in-person meetings (we had one secret exchanged by courier later), the final exchange could be done securely online. Every part of the secret had to be exchanged independently of any other part to maintain the level of security provided by OTP.

To make the three RCauth instances appear as a single instance, the three sites set up HAproxies to as HTTP load-balancing proxies in front of the delegation servers as depicted in Figure 5.1. These HAproxies each announce the same IP address (both IPv6 and legacy IPv4) using BGP Anycast. This achieves rapid (seconds-scale) failover without operator intervention (at the time of writing, anycast is supported by Nikhef and GRNET; STFC's data centre that hosts RCauth is going through a series of network reconfigurations in the second half of 2022 and should also eventually allow anycast). We considered other solutions, e.g., based on DNS with very

short Time to Live (TTL), but in the end chose the anycast solution for its speed, reliability, and ease of configuration.

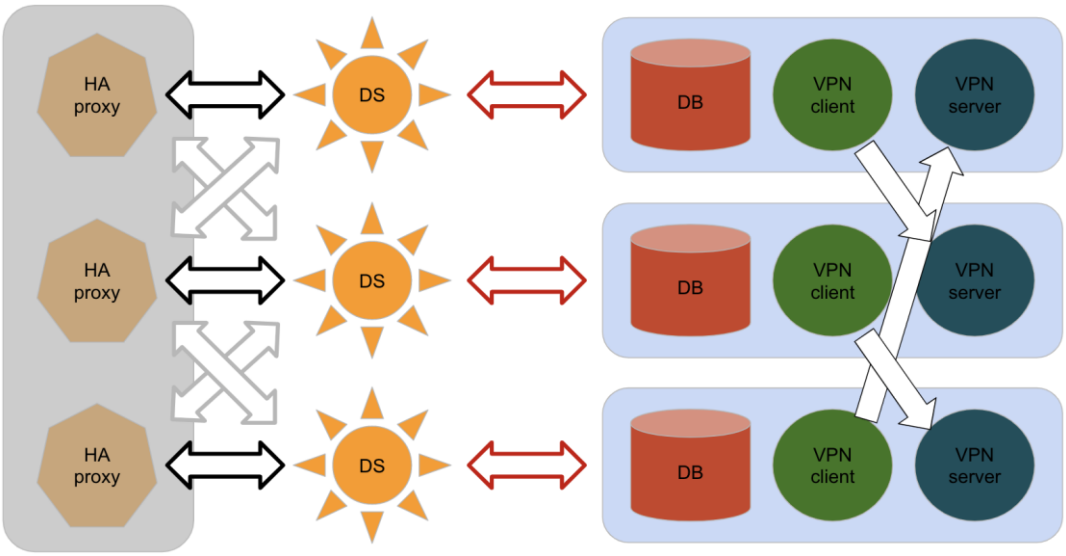

Figure 5.1: High availability deployment architecture of RCauth CA Delegation Server

The same anycasted-HAproxies are also used to proxy and load-balance the HTTP traffic towards the WAYF servers running at GRNET and Nikhef, such that the entire RCauth set of servers is now fully high-available (see Figure 5.2). Re-using the same HAproxy for both the DS and WAYF makes the maintenance effort smaller but effectively without impacting the high availability.

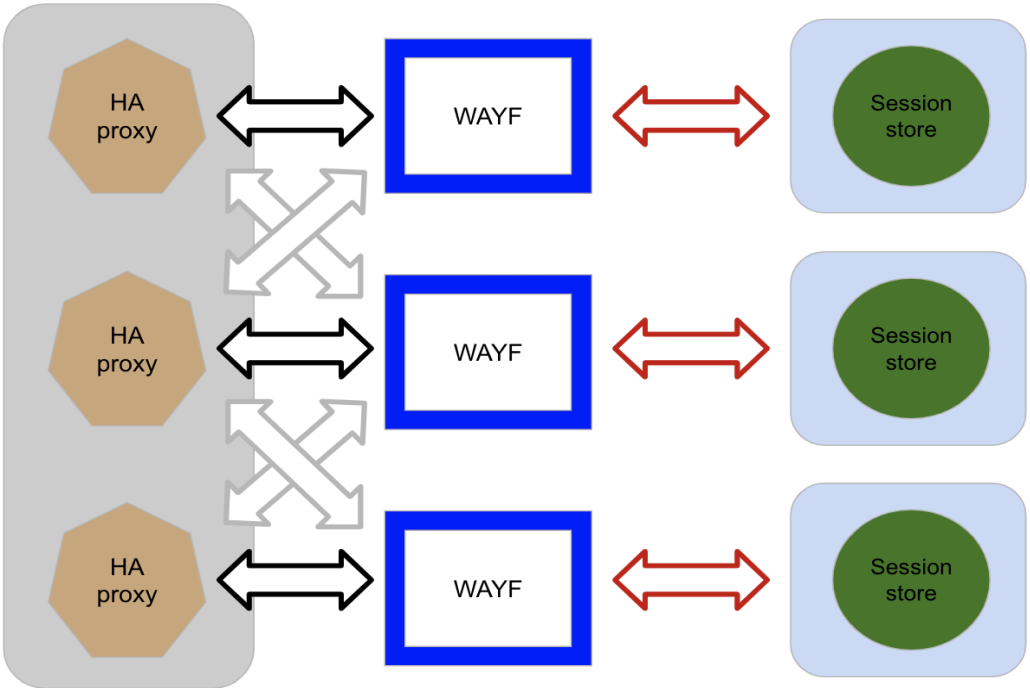

Figure 5.2: High availability deployment architecture of RCauth CA WAYF

The activities during the reporting period also included regular tasks:

- **Operations** - weekly operations calls, to track activities that require coordination between the sites.
  - These are documented in the internal wiki and summarised in the monthly WP7.3 calls
- **Maintenance** of systems - databases, VPN (virtual private network), HAproxy etc.
  - This includes applying (security) updates of the different servers.

  - There are minor service components, mostly custom-made shell scripts, which support the communications between the main parts of the system. These also need configuring, maintenance and documenting.
- **Ensuring the EUGridPMA is kept informed of the progress**.
- **Dissemination of activities**. While RCauth itself is a specific highly available service, there is a lot of interest in reusing the technologies developed for RCauth to make other web services highly available.

The need to keep EUGridPMA informed arises because the IGTF accreditation of RCauth is established through the EUGridPMA. It is therefore important that the EUGridPMA approves any RCauth plans that would significantly impact the overall architecture prior to their implementation and is kept informed of the experiences and lessons learned from the process.

6 IGTF Certificate Proxy

The IGTF Certificate Proxy is an Identity Provider (IdP) which supports authentication via X.509v3 certificates. It is configured to accept only certificates issued by Certification Authorities accredited under the IGTF Classic or MICS profiles. The IdP has been registered in eduGAIN and is available as an authentication option for the EOSC-Core Infrastructure Proxy. Thus, users who want to use their IGTF Certificate for accessing a service can select the "IGTF Certificate Proxy" in the Identity Provider discovery page. They are then redirected to the IGTF Certificate Proxy IdP, which requires certificate authentication. During the authentication process, the IGTF Certificate Proxy IdP extracts the user information that is encoded in the certificate and makes it available to relying parties in the form of SAML assertions.

The operation of the IGTF Certificate Proxy required technological upgrades of the underlying framework and libraries to take advantage of new features and robustness, as well as automation of deployment tasks to ensure uninterruptible and performant operation. There is both a production and demo instance of the IGTF Certificate Proxy. During the reporting period, both instances were configured to operate in High Availability in Hot-Standby (Active-Passive) mode, i.e., there is a primary node that handling all of the traffic. The other node is ready and waiting to take over should the primary fails.

Furthermore, the IGTF Certificate Proxy was configured to release certificate information according to the voPerson v2.0 specification, i.e. using the following SAML attributes:

- Distinguished Name of X.509 Certificate Subject: voPersonCertificateDN (object identifier: 1.3.6.1.4.1.25178.4.1.3)
- Distinguished Name of X.509 Certificate Issuer: voPersonCertificateIssuerDN (object identifier: 1.3.6.1.4.1.25178.4.1.4).

The number of monthly user logins through the Certificate Proxy is illustrated in Figure 6.1.

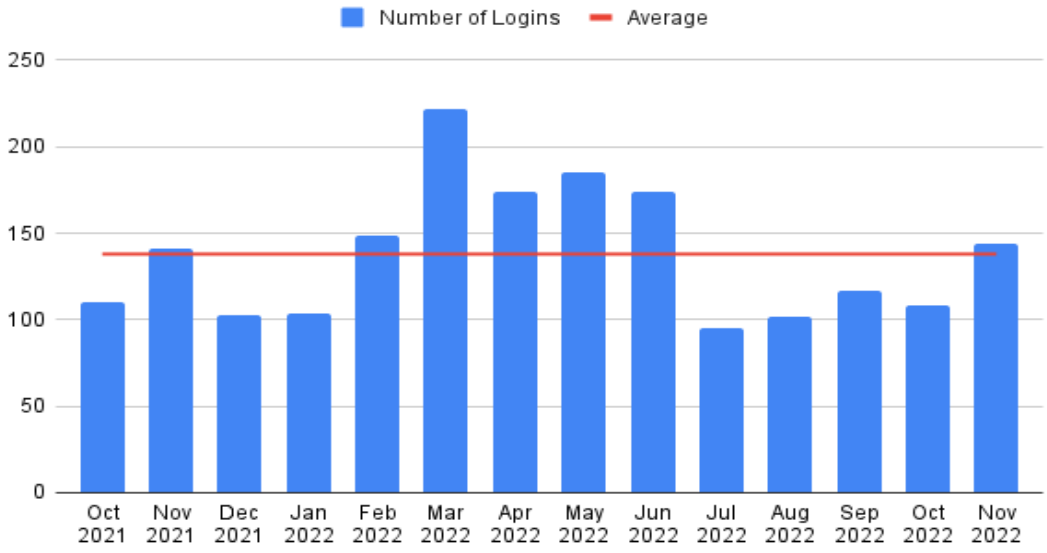

Figure 6.1: Monthly number of user logins through the IGTF Certificate Proxy

The cumulative sum of unique users accessing services through the IGTF Certificate Proxy is illustrated in Figure 6.2.

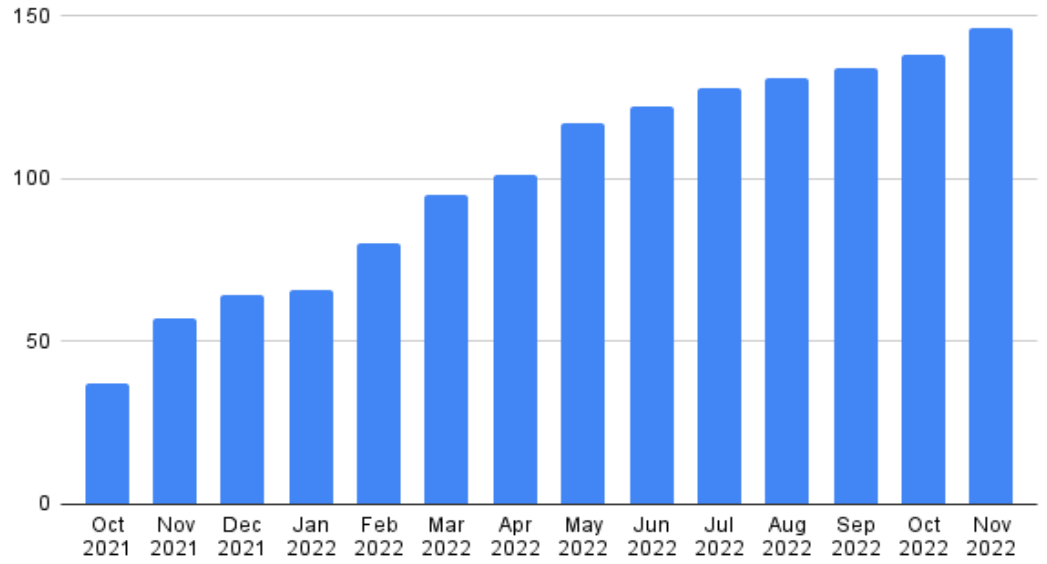

Figure 6.2: Cumulative sum of unique IGTF Certificate Proxy users per month

7 EOSC-Core AAI Fabric Monitoring

The EOSC-Core AAI Fabric Monitoring is a monitoring system for checking the state of the EOSC-Core AAI services. This is not meant to replace the central monitoring service from Task 4, but rather to provide operational monitoring for the AAI services.

During the reporting period, the EOSC-Core AAI Fabric Monitoring was deployed and configured to monitor the following AAI services:

- EOSC-Core Infrastructure Proxy (production and demo environment)
- IGTF Certificate Proxy (production and demo environment)
- RCauth Online CA (production and acceptance environment)

The services have been organised into three distinct host groups as shown in Figure 7.1.

	Host Group	Host States		Service States	
2	igtf-proxy	2		22	22
12	infra-proxy	12		108	108
16	rcauth	3	13	17	17

Figure 7.1: Host group configuration in EOSC AAI Fabric Monitoring

The monitoring checks include:

- IPv4/IPv6 network connectivity,
- network response time,
- DNS,
- validity of host certificates,
- validity of signing/encryption certificates used in SAML 2.0 metadata

Furthermore, monitoring agents were deployed on the hosts of the EOSC-Core Infrastructure Proxy and the IGTF Certificate Proxy to allow tracking CPU utilisation, number of processes, as well as memory and disk usage.

The EOSC-Core AAI Fabric Monitoring service provides a dashboard allowing members of the operations teams to get an overview of the state of the hosts and the configured service checks. The summary page for the monitored hosts and services is shown in Figure 7.2.

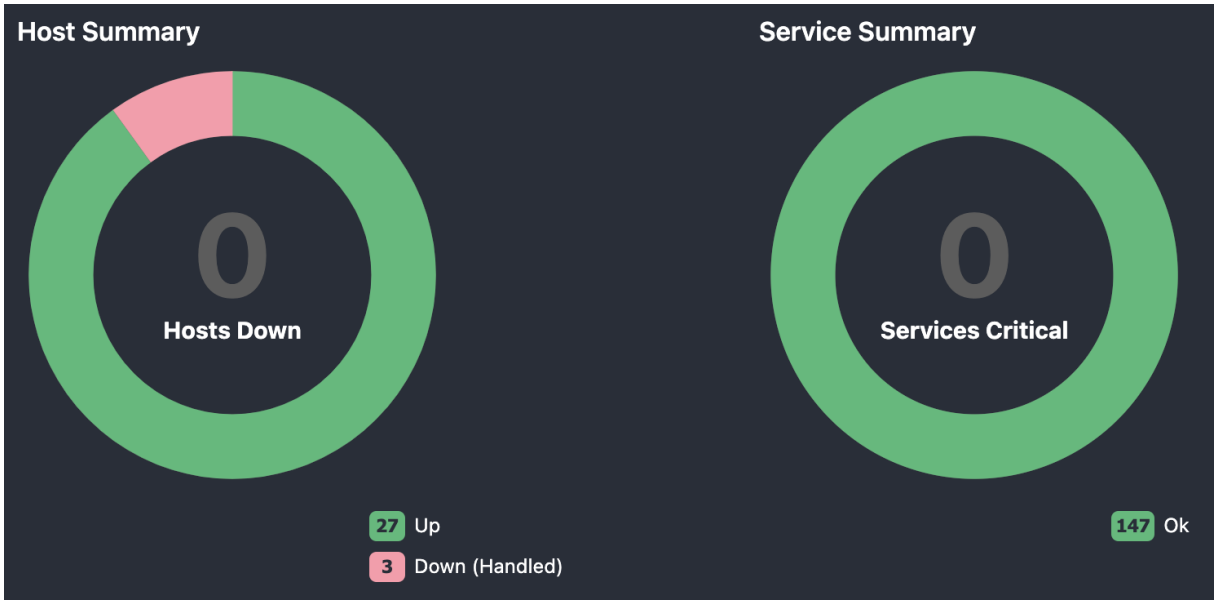

Figure 7.2: EOSC AAI Fabric Monitoring dashboard: Tactical Overview

It is worth noting that the EOSC-Core AAI Fabric Monitoring has been configured to send notifications to operation teams whenever important conditions are identified. At the same time, members of the operations teams can view the monitoring state through a web dashboard.

8 The EOSC AAI Federation

The AAI task in EOSC Future started the technical implementation of the EOSC AAI Federation based on the technical and policy guidelines described in the EOSC Authentication and Authorisation Infrastructure report published by the European Commission, Directorate-General for Research and Innovation in January 2021 [EOSC-AAI-2021].

Although the EOSC AAI Federation policies and practices are technologically agnostic, support for multilateral federation is still under development for OpenID Connect in the OpenID Forum. Thus, the initial implementation of the EOSC AAI Federation relies on SAML2 for the federation aspects. This means that services or authentication providers, which will be using OpenID Connect, will have to use the OpenID Connect to SAML2 proxy bridge provided by the EOSC AAI Federation to participate.

As shown in the Figure 8.1, members of the EOSC AAI Federation include organisations operating Community AAIs and/or Infrastructure Proxies for international access and organisations providing services in EOSC.

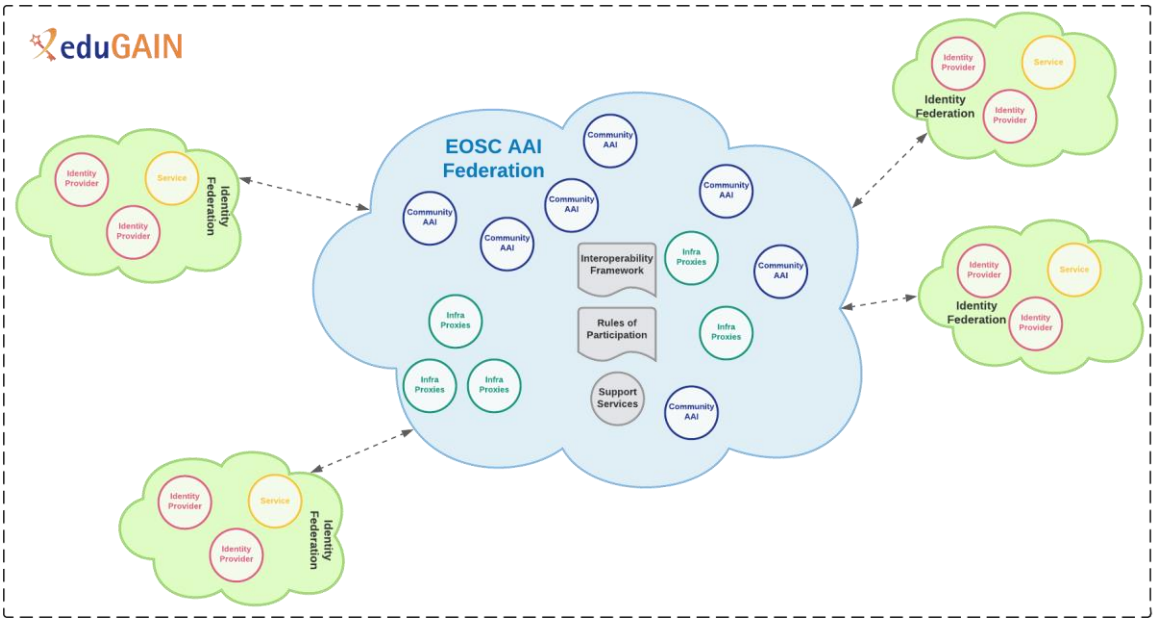

Figure 8.1: EOSC AAI Federation

Members can register Community AAI and Infrastructure Proxy entities in the EOSC AAI Federation Registry and have their metadata published in the EOSC AAI Federation metadata. Members may also be members of a peer federation of the EOSC AAI Federation via eduGAIN. In such a case, they shall not register any entities that are already registered in a Peer Federation and their entities will be imported via eduGAIN. Members providing end services to EOSC shall register their entities in the national federation of the country they are operating from, and they will be imported via eduGAIN or to an Infrastructure Proxy connected in the EOSC AAI Federation.

In the reporting period, the focus for the EOSC AAI Federation was to implement the supporting services required for the operation of the federation and to connect the operational AAI services (Community AAIs and Infrastructure Proxies) of the Cluster Research Infrastructures.

The technical infrastructure is currently comprised of the Metadata Registry and the OpenID Connect Connector.

The Metadata Registry is used to manage the federation entities and their metadata. In the reporting period, the Metadata Registry holds information about the following entities:

- Life Science Login Service. The Life Science Login service is used by EOSC Life, the Life Science Cluster Research Infrastructure, which brings together the 13 Life Science 'ESFRI' RIs. The Metadata Registry has information about the Community AAI component of the Life Science Login Service and EOSC Life's Infrastructure Proxy.
- UmbrellaID AAI Service. The UmbrellaID AAI Service is used by PANOSC and ExPaNDS, which bring together 6 European PAN Research Infrastructures and 10 European national PAN RIs in one Cluster RI. The Metadata Registry has information about the Community AAI component of the UmbrellaID AAI Service and the corresponding Infrastructure Proxy of this community.

- ESCAPE IAM Service. The ESCAPE IAM Service is used by the ESCAPE Science Cluster, the European Science Cluster of Astronomy & Particle Physics ESFRI Research Infrastructures. The ESCAPE IAM Service is using the OpenID Connect Connector of the EOSC AAI Federation. The Metadata Registry has information about the Community AAI component of the ESCAPE IAM Service and their Infrastructure Proxy.
- CESSDA AAI & DARIAH AAI. The CESSDA AAI is used by the CESSDA ERIC, and the DARIAH AAI is used by DARIAH in the SSHOC Science Cluster. The DARIAH AAI provides catch-all AAI services for the rest of the SSHOC RIs that do not have an operational AAI. The Metadata Registry has information about the Community AAI components of the CESSDA AAI and DARIAH AAI and their respective Infrastructure Proxies.
- EOSC-Core Infrastructure Proxy connects the EOSC-Core Services developed and operated in the project. The Metadata Registry has information about the Infrastructure Proxy component of the EOSC-Core Infrastructure, but no Community AAI functionality is required.
- eduTEAMS. The eduTEAMS Service is provided by GÉANT to small and medium sized communities who want to get started with their virtual collaborations and take full advantage of the federated access without having to deal with the complexity of operating and supporting their own AAI. eduTEAMS supports multiple communities on the same platform and provides everything required to securely collaborate and use services available both to the GÉANT community and through the European Open Science Cloud. Service Providers from small and medium sized communities can use the eduTEAMS Infrastructure Proxy to connect their services to the EOSC AAI. The Metadata Registry has information about the Community AAI component of the eduTEAMS Service and its Infrastructure Proxy.
- GÉANT AAI Service. The GÉANT AAI Service is provided to the GÉANT Community to access services and share resources provided by the GÉANT and the European Open Science Cloud and to Service Providers who want to make their services available to the GÉANT Community and EOSC. The Metadata Registry has information about the Community AAI component of the GÉANT AAI Service and the Infrastructure Proxy
- Check-in. The Check-In service is provided by EGI and it is a proxy service that operates as a central hub to connect federated Identity Providers (IdPs) with EGI service providers. The Metadata Registry has information about the Community AAI component of the Check-in Service and its Infrastructure Proxy.
- B2ACCESS. The B2ACCESS service is provided by EUDAT, and it is a federated cross-infrastructure authorisation and authentication proxy for user identification and community-defined access control enforcement. It allows users to authenticate themselves using a variety of credentials providing federated access and single-sign-on to services and service providers in a trusted way. B2ACCESS offers communities and service providers an AARC compliant AAI proxy ready to be integrated within the EOSC AAI Federation. The Metadata Registry has information about the Community AAI component of the B2ACCESS Service and its Infrastructure Proxy.
- OpenAIRE AAI. The OpenAIRE AAI is provided by OpenAIRE and enables researchers to securely access and share common resources and services. The Metadata Registry has information about the Community AAI component of the OpenAIRE AAI Service and its Infrastructure Proxy.

In addition, the EOSC AAI Federation Metadata Registry automatically includes Identity Providers from eduGAIN that comply with the Security Incident Response Trust Framework for Federated Identity (Sirtfi). At the time of this writing **1249 Identity Providers are automatically imported via eduGAIN**.

The Metadata Registry has implemented preliminary support for AARC-G079 "AARC Community-based Access Entity Category". This guideline document is not part of the AARC Blueprint Architecture, as it has not been approved by AEGIS yet. This is expected to happen in Q1 of 2023.

The OIDC Connector is a component provided by the EOSC AAI Federation, which allows Community AAIs that support only OpenID Connect towards services, to still be able to connect to the EOSC AAI Federation.

For each Community AAI that uses the OIDC Connector, it creates one virtual OpenID Connect Client interface and one virtual SAML2 Identity Provider Interface. The SAML2 Identity Provider Interface presents all the

metadata information of the Community AAI as if it were an interface of that Community AAI and it is registered automatically in the EOSC AAI Federation.

Whenever an authentication request is received by the virtual SAML2 IdP Interface of the OIDC Connector, it is translated to an OpenID Connect Authentication Request towards the respective OIDC Provider interface of the Community AAI. The authentication response from the OIDC Provider interface of the Community AAI is received by the virtual OIDC Client interface of the OIDC Connector and then translated to a SAML2 Authentication Response.

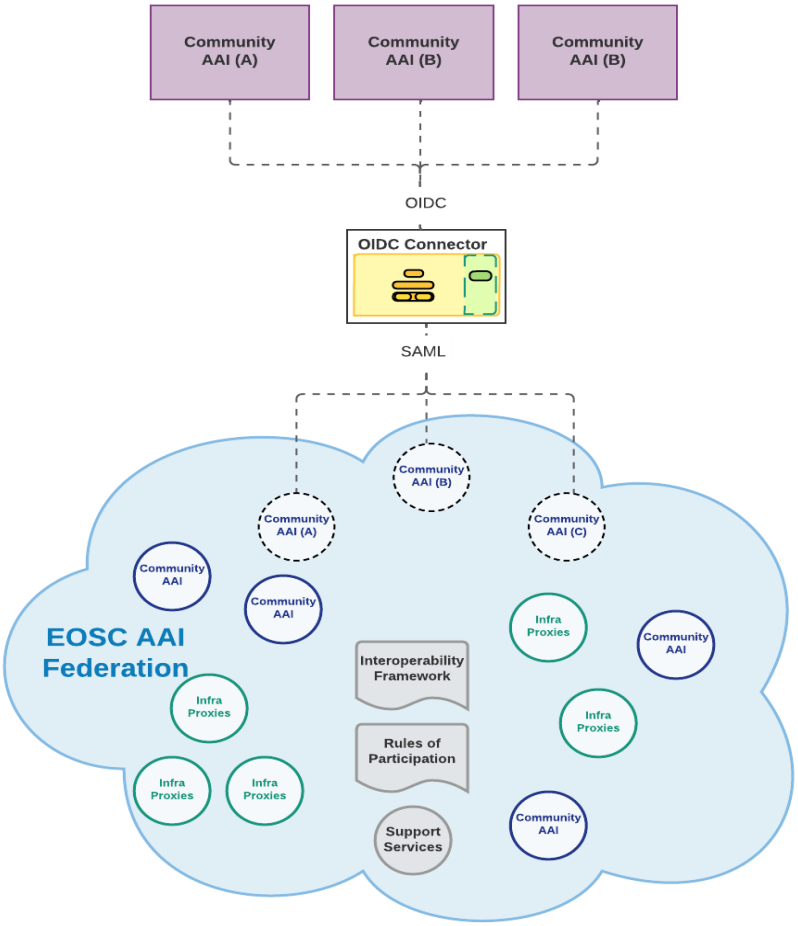

Figure 8.2: EOSC AAI Federation OIDC Connector

The EOSC AAI Federation supporting services are cloud-native components running on the top of the GÉANT Core AAI Platform on Amazon AWS (in the Frankfurt region) procured using the OCRE Framework. The GÉANT Core AAI platform implements Active-Active High Availability Configuration and can scale horizontally along with the demand.

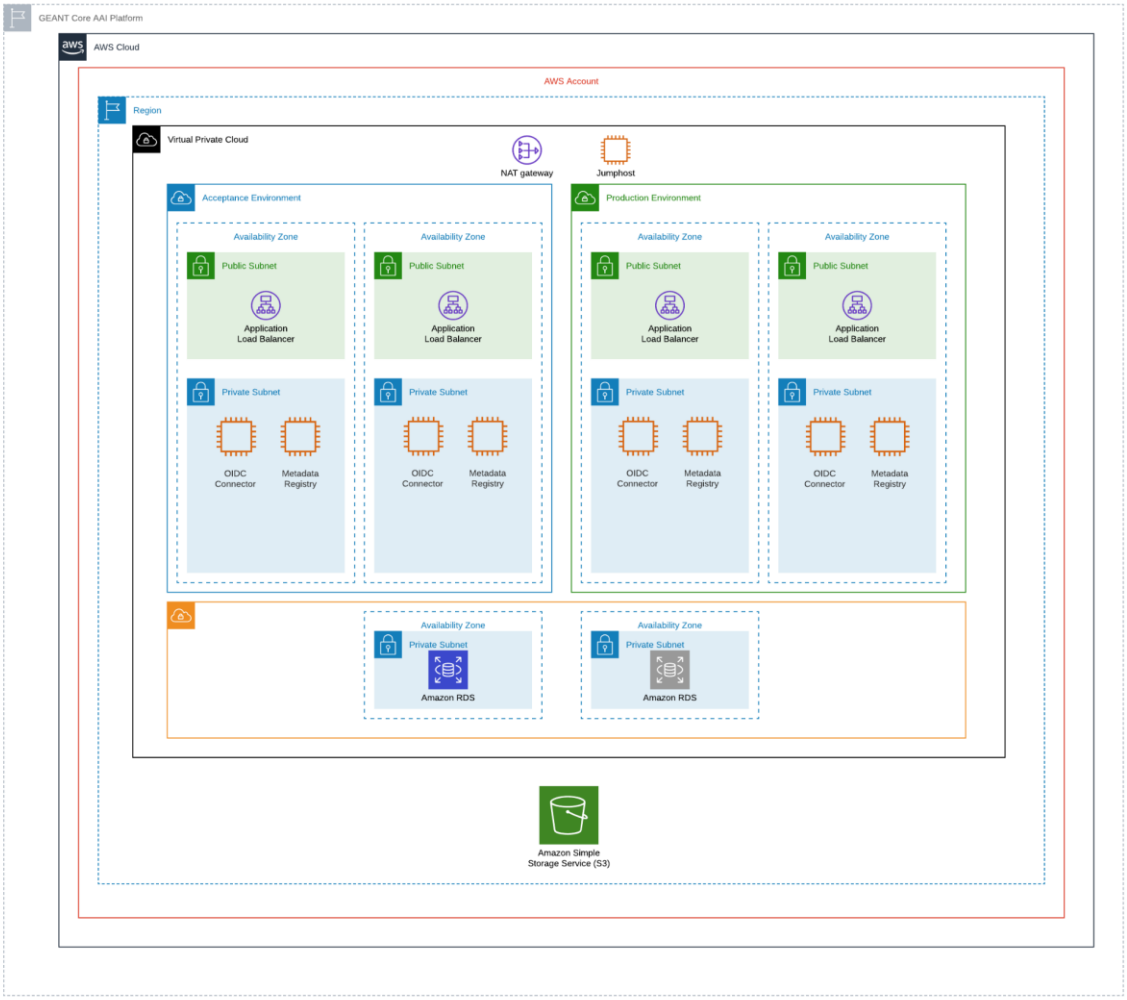

Figure 8.3: AWS Infrastructure

9 Work with the Science Clusters, Research, and e-Infrastructure Providers

To connect the operational AAI services (Community AAIs and Infrastructure Proxies) of the Cluster Research Infrastructures (described in Section 8), the AAI task organised regular coordination calls with the Science Clusters, the Research Infrastructures, and the e-Infrastructure Providers. In the reporting period, twelve (12) coordination calls took place.

The initial goal of these support activities was to benchmark the AAI readiness of the Science Clusters and the e-Infrastructures against the policy and technical requirements for connecting to the EOSC AAI Federation. As part of this process, the AAI support team prepared a questionnaire to gather initial input from each Science Cluster and then based on this input, the support team followed up with one-to-one calls with each of the Science Clusters to prepare cluster specific action plans.

The benchmarking showed that the e-Infrastructures, PANOSC and EOSC Life Science Clusters and the CESSDA ERIC had the most mature AAI implementations and would be the first to be added in the EOSC AAI Federation in April 2022.

ESCAPE was the third Science Cluster, which was equally mature in terms of functionality, but which required the availability of the OIDC Connector to join the EOSC AAI Federation. The OIDC Connector was made available in May 2022 and then we were able to proceed with the integration of ESCAPE in the EOSC AAI Federation.

The DARIAH AAI was integrated in the Summer of 2022 and in addition it was agreed that it would act as a catch-all solution for the SSHOC cluster, supporting users and service providers from the SSHOC RIs that do not have access to an operational AAI.

The CLARIN AAI, even though it is compatible with the AARC BPA, does not implement proxied solutions. Instead, it is a federation of SAML2 service providers and will follow the generic service provider onboarding through the national academic federations. In the future, if CLARIN needs to connect service providers using OIDC, it may use one of the available Infrastructure Proxies or use the DARIAH AAI. The option of implementing their own Infrastructure Proxy is also a possibility.

For the ENVRI-FAIR Science Cluster the situation was different as, in parallel with the EOSC Future activity, the Cluster in its own implementation project had an internal AAI task force whose primary responsibility was to evaluate the AAI options for ENVRI-FAIR. This work concluded that in ENVRI-FAIR there would not be a cluster-wide AAI, and instead each Research Infrastructure would implement their own AAI solution. LifeWatch started the implementation of its own AAI in October 2022 with the support of the EOSC Future AAI team. In January 2023, the ENVRI-FAIR Cluster organised a workshop on the topic of the AAI where the EOSC Future support explained the requirements and expectations of the EOSC AAI and agreed on a course of action for delivering AAI capabilities to ENVRI FAIR with an initial focus to the Science Projects that need to use the EOSC services. This activity is ongoing, and we expect to have the first results before the beginning of Summer 2023.

10 The work ahead

10.1 International Standardization and Harmonization activities

- OAuth 2.0 Proxied Token Introspection (AARC-G052)
- Establishing trust between OAuth 2.0 Authorization Servers (AARC-G058)

10.2 RCauth Certification Authority

The activities planned for the remainder of the project include the following:

- STFC's data centre will get fully refreshed hardware - new servers, new Hardware Security Modules (HSMs) in Q1 2023. The upgrade of a router "closer" to the RCauth infrastructure is also planned which, combined with the network restructuring, should enable anycast.
- Deploy a WAYF component of RCauth at STFC (depends on anycast). This will be fronted by the anycasted HAProxies (see Section 4.2), as is currently the case for the WAYF components provided by Nikhef and GRNET.
- Configure the acceptance instances of RCauth running on the three sites so that they appear as a single instance. The plan is to use a DNS based solution rather than BGP anycast. While not as robust as anycast, a DNS solution should be sufficient for the availability/reliability requirements of the acceptance environment which is used for testing changes (e.g., upgrades of service components or the integration of new clients/master portals) before moving to production.
- Further dissemination activities. While RCauth itself is a specific highly available service, there is a lot of interest in reusing the technologies developed for RCauth to make other web services highly available. We have planned to write up a paper for peer reviewed publication, and plan to make it easier for others to reuse parts of RCauth.

10.3 EOSC-Core Infrastructure Proxy

The following activities are planned:

- Connect additional EOSC-Core services to the production instance of the EOSC-Core Infrastructure Proxy. At the time of writing, the following services have been identified:
  - EOSC Accounting for Services (connection is currently being tested in the acceptance environment)
  - EOSC Knowledge Hub

- Complete migration of services to the new OpenID Connect/OAuth 2.0 Provider (OP) component of the EOSC-Core Infrastructure in production. As already stated in Section 5.2, the new OP component improves compliance with the OpenID Connect specification and OAuth 2.0 Best Current Practices and adds support for the OAuth 2.0 Proxied Token Introspection (AARC-G052). The latter will allow EOSC resource servers (e.g., EOSC Resource Catalogue API) to handle tokens which are issued by other trusted OAuth 2.0 Authorisation Providers within the EOSC AAI Federation.
- Add support for inferring and constructing affiliation information (voPersonExternalAffiliation) according to AARC-G057.
- Investigate support for mechanisms allowing the dynamic establishment of trust with other OAuth 2.0 Authorization Servers within the EOSC AAI Federation.

10.4 IGTF Certificate Proxy

The service roadmap for the IGTF Certificate Proxy includes adding support for expressing user identifiers using the `voPersonID` SAML attribute, as per AARC-G026.

10.5 EOSC-Core AAI Fabric Monitoring

No further activities are planned, apart from the regular maintenance of the underlying infrastructure and software components.

10.6 EOSC AAI Federation

The following activities are planned:

- Implementation of a one stop shop for Organizations and Community AAIs to register in the EOSC AAI Federation.
- Initial implementation of the AARC-G058 specification: "Establishing trust between OAuth 2.0 Authorization Servers". This specification is still at the early stages of the standardisation process in the AARC Community and AEGIS.
- Application to register as a relying federation in eduGAIN
- Ability to login with national eIDs via eIDAS

10.7 Work with the Science Cluster, Research, and e-Infrastructure Providers

The following activities are planned:

- Continue the work with ENVRI-FAIR to ensure that all the ENVRI-FAIR Science Projects will be able to use the EOSC services
- Continue the work with the Cluster RIs, Research and e-Infrastructure Providers
- Liaise with the INFRA-07 project to onboard AAI services implemented in those activities.

11 Conclusions

The AAI task in the project delivered the first implementation of the EOSC AAI Federation, enabling the Research Infrastructures in the Science Clusters and the e-Infrastructures to fully utilise the capabilities of the EOSC AAI. By M20, all the operational AAI Services in the Science Clusters and the e-Infrastructures, including the EOSC-Core Infrastructure Proxy, have been added to the EOSC AAI Federation.

The AAI team developed new operational procedures along with a clear policy and technical requirements for connecting end services to the EOSC-Core Infrastructure Proxy, turning this component from an integration tool to an operational production service, key for the delivery of the EOSC-Core Services.

To achieve these goals, the AAI team had to actively engage with the international standardisation community and lead the development of six new AARC interoperability guidelines. Four have already become part of the AARC Blueprint Architecture. The other two are expected to be finalised in Q1 of 2023. In addition, the team has led the development of the next version of the EOSC AAI Architecture, scheduled to be published in January 2023 by the EOSC Association.

All this work would not have been possible without the strong collaboration between the AAI team and the Science Clusters, the Research Infrastructures, and the e-Infrastructures.

In the following months until the end of the project, the focus turns to widen the reach of the EOSC AAI, with more communities and services joining from the EOSC 07 projects and providing onboarding paths for commercial services and services from other sectors.