10-05-2019

Deliverable DSA1.5: How-to Deploy Pilot Results

Deliverable DSA1.5

Contractual Date:	30-04-2019
Actual Date:	28-05-2019
Grant Agreement No.:	730941
Work Package:	SA1
Task Item:	TSA1.4
Lead Partner:	RETI
Document Code:	DSA1.5

Authors: Simone Visconti; Arnout Terpstra; Mario Reale; Ioannis Igoumenos

Abstract

This document is the final deliverable of Service Activity 1 (SA1) within the AARC2 project for the implementation of pilots with the user communities. It reports on the overall outcome of the pilots carried out within AARC, highlighting how their implementation matches the AARC blueprint architecture and its identified functional components. Results are reported for each community including initial requirements, a description of the implemented system and an overall assessment of the outcomes with respect to the original goals set.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 730941 (AARC2).

Table of Contents

Table of Figures

Table of Tables

1 Introduction

Service Activity 1 Pilots (SA1) have demonstrated the feasibility of deploying Authentication and Authorisation Infrastructures (AAI) for research communities and e-infrastructures that fit the overarching AAI model defined by the AARC Blueprint Architecture (BPA) [AARC BPA]. To this end, the activity demonstrated through (pre-)production pilots that:

- The AARC BPA can be instantiated to fit research communities' requirements and deployed and operated in production environments.
- Communities are enabled to design and choose an e-infrastructure provider (or more) that can deliver AAI services compliant with the AARC BPA or operate the AAI themselves.
- User/group information can be retrieved from distributed group managements and attribute providers. This information in combination with the affiliation that is provided by the user Identity Provider is used for authorisation purposes.
- Communities at an early, initial phase in the design of their AAI solution benefit the most from the support of exports. Targeted work aimed at supporting them in the design of a scalable, robust, AARC BPA-compliant architecture has proven to be fundamental towards avoiding issues and bottlenecks in the implementation of the AAI component.
- All relevant issues for user communities have been addressed by AARC, so that in all pilots carried out the fundamental community requirements have successfully been fulfilled and the VOs have been able to take over from the project to pursue their sustainable solution to AAI in the future.

A total of nine research community pilots have been deployed under the umbrella of SA1. This document introduces an overview of the pilot approach used in SA1, then proceeds with a description of each pilot and the requirements it covers.

2 Pilot Approach

When planning the pilot activities, the first fundamental step has been to gather the requirements owned by each user community. The user communities showed different levels of maturity and skills in the AAI domain. Therefore, the task organised structured interviews to guide each community to express its requirements in a well-documented, consistent way.

The interviews touched on topics such as their current (initial) solution for AAI, the level of involvement of their users in eduGAIN, and the resolution in the authorisation model while accessing computing resources and data.

The subsequent phase involved the actual design of a scalable solution for the community, based on existing software products and tools, and framing its architecture within the context of the AARC BPA.

Interactive sessions and hands-on joint discussion led to drawing up the proposed architecture on whiteboards and in reference documents (e.g. in the AARC2 wiki and other deliverables).

Implementation was then started in a joint effort by the AARC SA1 team and the community AAI engineers.

The implementation phase lasted for some weeks, even months in some cases, since the solutions had to be made to fit the current community AAI solution. In some cases, SA1 provided effort to design fully automated deployable suites, aimed at easing the deployment of the pilot infrastructure.

Overall, in order to progress from initial research community requirements to a fully functional implementation of an AAI according to the AARC BPA, the process illustrated below was used within SA1:

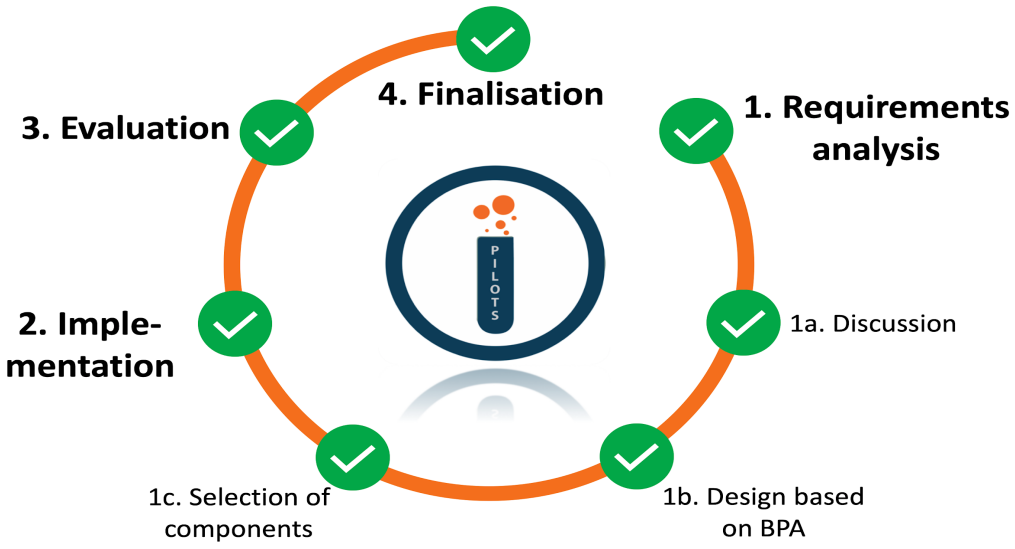

Figure 1: Pilot approach

2.1 Requirements Analysis

2.1.1 Discussion

Research communities should specify their requirements for an AAI as the first step in the process. Not only does this help the research community to define the end goal(s), it also helps it in prioritising which aspects of an AAI are more important and should therefore be deployed first and which can be introduced in a later iteration. For an example see the requirements specified by the Life Sciencescommunity [LSrequirements] or use the FIM4R v2 paper as guidance [FIM4Rv2].

2.1.2 Design Based on BPA

Once the requirements are clear, a first design based on the AARC BPA can be drawn up. Having AAI expertise available during this step is crucial. During AARC2, this expertise was provided by the (non-research community) partners who also participated in the project. Most of the expertise was provided by National Research & Education Networks (NRENs), who are generally responsible for providing a national AAI infrastructure for research and education in their respective countries.

2.1.3 Selection of Components

Based on the design drawn up, components should next be selected (i.e. which software products are going to be used). It is important to already think carefully about who is going to maintain the production AAI service once it is ready; introducing new technologies with a steep learning curve could make maintaining the infrastructure more difficult. It should also be considered that e-infrastructure providers might already offer certain components 'as-a-service'. This could result in the research community's AAI being significantly easier to build and maintain.

2.2 Implementation

After the initial design and selection of components are completed, the implementation phase begins. During this phase, all components are rolled out and integrated with each other, and possibly with existing systems already part of the research community's infrastructure.

2.3 Evaluation

During the implementation phase, developers should use test tools (e.g. the AARC DiY IdP) while setting up the individual components. Once all BPA components are in place and configured, the evaluation phase can start. During this phase, typically some research community's services are connected to verify whether the AAI functions as expected.

2.4 Finalisation

Once the solution is working in a (pre-)production setting, existing users and services can be migrated to the new AAI. Technical documentation (internal) should also be written; this will be essential for the community service operators to connect their service to the newly built AAI.

3 Pilot Results

This chapter reports on the results of each pilot carried out as part of the AARC2 project. The FIM4Rv2 set of requirements were initially used to drive the discussion with each community and to scope the content of the pilots. However, as research communities involved in the pilots gained a better knowledge about the AARC BPA, they started to scope their own requirements.

3.1 Cherenkov Telescope Array (CTA)

CTA is a community of astrophysics users which already had its own AAI solution in place, and represents for AARC, in this respect, a very good example of how to address the needs of a community who has already developed an AAI. In this case their AAI solution was based on a SAML stand-alone, catch-all Identity Provider, integrated with a Group management tool used for Authorisation on selected service providers. The CTA pilot provides a non-invasive solution to simplify access to CTA services from eduGAIN and the CTA community. The work which has been carried out in the CTA pilot of AARC is aimed at providing the CTA community with eduGAIN authentication services ensuring at the same time a way to onboard this scientific community into eduGAIN. An infrastructure has been deployed based on the model proposed by the AARC Blueprint Architecture to enable the management of users coming from both eduGAIN Identity Providers and the CTA standalone IdP. The core component of the new infrastructure is the SATOSA IdP/SP proxy, as the central AAI layer to serve the CTA community of users. In addition to that, an external attribute authority [Comanage] has been plugged into the proxy in order to manage the user enrolment process, ensure injection of additional user authorisation attributes and allow for account linking whenever appropriate as requested by the users and granted by the manager of the collaboration. The AARC CTA pilot system has been successfully tested by the CTA AAI experts which have been able to authenticate and get authorised on specific CTA service providers.

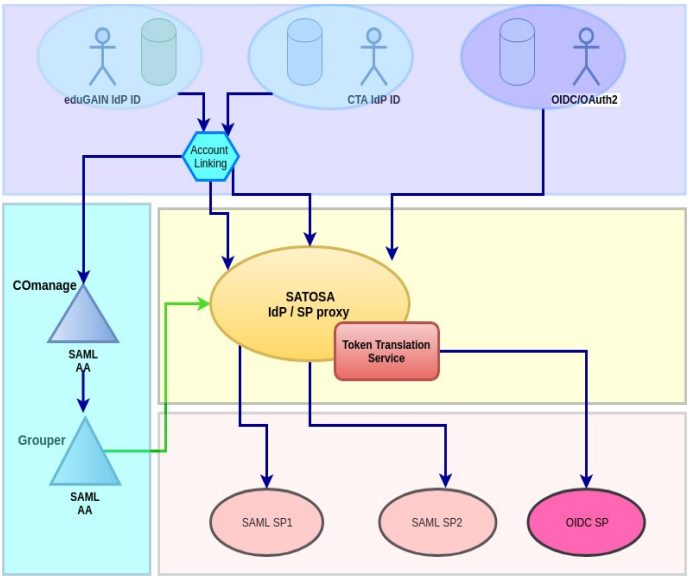

Figure 2: Architecture based on BPA for CTA

The designed workflows, supported by the SATOSA proxy and its implemented microservices, have been proven to work and be reliable and to support the desired authentication and authorisation processes.

3.1.1 Technical description

The core component of the pilot is the SaToSa IdP/SP proxy, which acts as main integrator of user identity in order to enable proper access to CTA services. SaToSa has been chosen given his feature of providing microservices on board, at the proxy level, and also facing IdP and SP sides. Microservices have been implemented to fetch key user information from the Comanage [Comanage] Attribute Authority and Grouper [Grouper] (which was already in use within CTA). The fundamental ePUID identifiers is used as primary identifier for the user and it is generated on board of the proxy. The user enrolment process goes via COmanage and queries are issued to link ePPN based IDs in COmanage to the generated ePUID. Grouper is subsequently used to extract the isMemberOf attribute to enable proper authorization on legacy CTA SPs.

Overall, both a complete user enrolment and management process has been designed, Identities from the CTA catch-all IdP and eduGAIN one's are linked, and both CTA legacy and new SPs are accessible making use of ePUID when possible.

3.1.2 Pilot Results

The pilot was successfully completed February 2019 and a pre-production infrastructure has been correspondingly set up in Catania at INAF. Given the positive result of the pilot, CTA is evaluating the possibility of moving it from the experimental phase to production, maintaining it and offering this service to the whole community. The AARC Blueprint Architecture was used as a model to design the pilot by clearly separating each component and its role in the system architecture. The pilot and its testbed will be maintained by INAF.

The main benefits for the CTA community can be summarised as follows:

- Successfully exploited an architecture capable of onboarding the whole CTA community to the eduGAIN trust model and flows.
- Include COmanage and Grouper as community tools to support attribute management and highly grained authorisation processes.
- Successfully integrated both legacy and new Service Providers of interest for the CTA community into the new AAI.
- Generation of the required ePUID as a unique, reliable identifier for the CTA users.
- Linking of identities between already existing CTA IDs and eduGAIN identifiers.

All the original goals of the pilot have been reached.

3.1.3 Requirements Table

Category	Requirement	How it was addressed	Achieved?
Identity lifecycle & linking	Account linking - The ability to link multiple identities together, whether held in parallel or succession. It depends strongly upon the release of an appropriately unique and persistent identifier.	Account linking is a feature provided by Comanage. It has been implemented by using ePUID as a strong unique and persistent identifier. This value is calculated in SaToSa and stored in COmanage.	✅
AuthZ	Single users 'management point	Users are managed by Grouper, where administrator can remove authorization for blocked users, authorise new users etc.	✅
	Deprovisioning – Meant as the ability to suspend or remove an individual's access when they no longer possess right of access.	Users can be deprovisioned by means of COmanage and Grouper.	✅
On-boarding & support	Proxy test/dev environment - a separate environment where new features could be tested separately from production	A development environment is set up and it is available to test new features and upgrades. It is separated from the production proxy.	✅

Table 1: CTA requirements

3.2 DARIAH

DARIAH is a European Research Infrastructure Consortium (ERIC), a pan-European infrastructure for arts and humanities scholars working with computational methods. It supports digital research as well as the teaching of digital research methods. It connects several hundred scholars and dozens of research facilities in currently 17 European countries, the DARIAH member countries. In addition, DARIAH has several cooperating partner institutions in countries that are not DARIAH members, as well as strong ties to many research projects across Europe. The people in DARIAH provide digital tools and share data as well as know-how. They organise learning opportunities for digital research methods, such as workshops and summer schools, and offer training materials for Digital Humanities.

The AARC2 DARIAH pilot consists of two individual goals:

1. Implementation of an SP-IdP proxy within the DARIAH AAI: according to the AARC BPA, communication between infrastructures should happen through dedicated infrastructure proxies. During this pilot, DARIAH implemented their own proxy solution based on Shibboleth. This proxy will be compliant with all relevant recommendations and guidelines developed within AARC. Therefore, this pilot can be seen as a real-world example of the architecture work carried out within AARC. As a side benefit, DARIAH-internal services will also be able to use this solution as it will move a lot of the complexity away from the individual services to the central proxy component.
2. Interoperability pilot between EGI and DARIAH. To showcase successful implementation of the DARIAH SP-IdP proxy, the second part of this pilot deals with interoperability between the DARIAH research infrastructure and the EGI e-infrastructure. The goal is to allow DARIAH users to transparently access EGI resources through EGI's own proxy solution (EGI CheckIn). As an initial use case, selected DARIAH users should be able to deploy and access virtual machines in the EGI infrastructure.

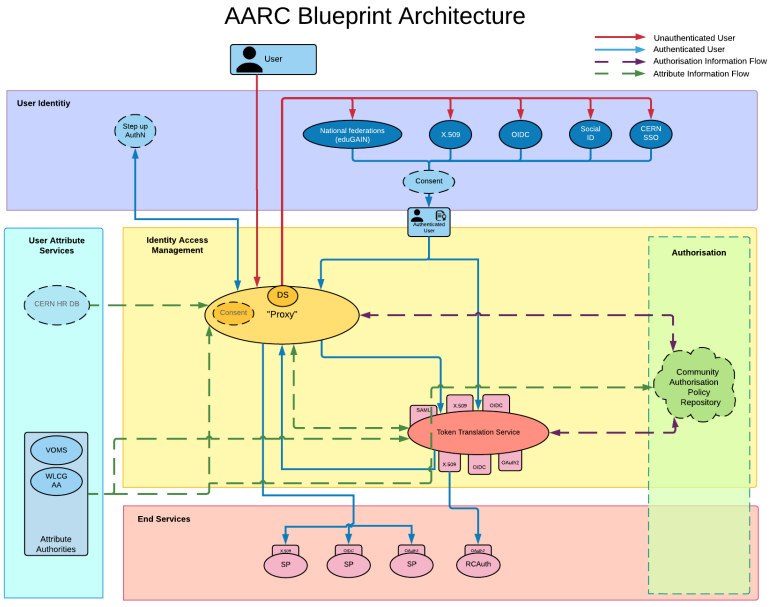

Figure 3: Architecture based on BPA for DARIAH

The DARIAH research infrastructure offers the DARIAH AAI as one of the core technical services for researchers in the arts and humanities. This enables researchers to log in to various services offered within DARIAH, by either using their own campus account or an account registered at the DARIAH homeless IdP. The DARIAH AAI adds information, such as group memberships specific to the DARIAH community or approval of general and optionally service specific terms of use, which can be used by services for authorisation decisions. Version 1 of the DARIAH AAI has been in production for several years and required every service to implement various details themselves, e.g. connection to eduGAIN, attribute query to DARIAH for the additional attributes, validation of policy attributes and blocking and redirecting the user to the DARIAH self-service portal if any of the information was missing or out of date.

In order to overcome these limitations and be in line with the AARC BPA, and therefore allow interoperability with other infrastructures, DARIAH decided to implement DARIAH AAI version 2. In this version 2, which was created as part of the pilot, a central proxy component, which follows the proxy layer of the AARC BPA, was added to complete the first goal (Implementation of an SP-IdP proxy within the DARIAH AAI) of the pilot. The implementation is based on the building blocks of the BPA and also follows various guidelines from the AARC project on how to convey information to other infrastructures. Examples include the unique identifier or how to express group memberships in order to allow DARIAH users access to EGI resources.

3.2.1 Technical description

The implementation of the proxy is based on Shibboleth IDP and SP. The decision to use this software was mostly done due to the prevalence of Shibboleth software in other components of the DARIAH AAI. As part of the pilot we've developed some plugins in order to connect Shibboleth SP and IDP and allow e.g. exchange of attributes and sessions between the two.

Within the DARIAH AAI, even before starting this pilot, a group and user management tool, called didmos LUI and developed by DAASI International, who runs the DARIAH AAI, was present. Part of the pilot was to extend Shibboleth IDP in a way, that it is capable to communicate with this tools, for example to check if a user is already registered or needs to visit the group and user management platform in order to register a DARIAH account.

Most of the other functionalities were implemented used exiting features of Shibboleth SP and IDP.

3.2.2 Pilot Results

The implementation of the proxy was completed in mid-2018 and since then all DARIAH services have been moved behind the proxy. Since the architecture was designed with backwards compatibility in mind, the transition process did not create any major issues. Using the proxy to connect to federated AAI is now much simpler for service operators, hence additional services have already been connected to the DARIAH AAI with the intention of continuing to add further services. In the second part of the pilot the DARIAH proxy was successfully connected to the development instance of EGI check-in. This includes attribute and entitlement mapping from DARIAH to EGI, as well as on-the-fly user provisioning within EGI. For this purpose, some plugins to the existing EGI check-in infrastructure were developed.

Within DARIAH the new proxy component is now part of the core AAI and will be operated as such in the future, including the enrolment workflows for DARIAH users on the EGI infrastructure. The DARIAH AAI is operated by partners in the DARIAH project and one of the main technical components used by researchers. There is a dedicated helpdesk available, which deals with both end user issues and also with service operators, who want to connect to the DARIAH AAI.

3.2.3 Requirements Table

Category	Requirement	How it was addressed	Achieved?
Discovery & usability	Smart discovery - IdP discovery should be "smart enough" to quickly and easily take a user to their appropriate home IdP, for example, show the user a short list tailored to them by home country, institute, e-Infrastructure, research community, project or other hints.	Discovery Service at the DARIAH proxy offers the DARIAH homeless IDP as an extra button, since most users actually use this IDP for authentication. Furthermore IDPs in eduGAIN, which have already been used for a DARIAH account, are placed on top of the list.	✅
AuthZ	Real-time authorisation – AuthZ decisions at an SP must be based on identity credentials and attributes; the assertions that validated have a short lifetime. Even within this short period it should be possible for the SP to look up real time status information, e.g. revocation lists and/or suspension lists.	Whenever users authenticate through the DARIAH Proxy their , the Proxy validates the account validity and the user attributes in real time, before allowing the user to access the service. Also, authZ information conveyed to services is generated at each login and reasonably up to date, since assertions have limited life-time.	✅
	User blocking - it must be possible for an Infrastructure or research community to block access to a service based on the presence of an identity credential in an operational suspension list or revocation list.	Users can be deactivated globally through the DARIAH user management portal.	✅

	Group management - Research communities must be able to add individuals to Groups, for use in AuthZ, Quota management and Accounting. Groups should be hierarchical and users can belong to more than one group.	The existing DARIAH group and user management, based on the didmos LUI software from DAASI International, who runs the DARIAH AAI, was integrated into the DARIAH proxy.	✅
Attribute release	Attribute management via the proxy	DARIAH adopted the AARC guidelines on identifier between infrastructures (G002) and also use such an identifier internally within DARIAH. Research services connected to the DARIAH proxy receive the full DARIAH attribute set, which includes R&S and the DARIAH Community Identifier.	✅
Security incident response	[Sirtfi] adoption - IdPs and SPs should meet the requirements of Sirtfi and assert this in metadata.	The requirements of SIRTFI are met by the DARIAH community and adding the SIRTFI entity category to eduGAIN metadata of the DARIAH Proxy SP is ongoing.	✅
Research community proxies	IdP/SP Proxies should be available via eduGAIN	The SP part of the DARIAH proxy is in eduGAIN.	✅
Security aspects related to the usage of the procy	Compliance with [SNCTFI]- The IdP-SP proxy puts requirements on the research communities service providers. Snctfi provides a framework to ensure all necessary elements are supported in order to preserve security and privacy.	DARIAH analysed what is required to comply with Snctfi requirements and were able to fulfil most of them. However, because there is not yet an AuP in place, not all Snctfi requirements can be considered fulfilled. This is work in progress and it will be re-evaluated in the future.	❌
Sustaining critical infrastructure	Sustainable operation of specified critical services	The proxy is now considered a core component in the DARIAH AAI and continued to be operated as such.	✅

Table 2: DARIAH requirements

3.3 EISCAT_3D

EISCAT_3D is a project to establish an international research infrastructure using radar observations and an incoherent scatter technique to conduct studies of the atmosphere and near-Earth space environment above the Fenno-Scandinavian Arctic, as well as to support the solar system and radio astronomy sciences. The radar system is designed to investigate how the Earth's atmosphere is coupled to space, but will also be suitable for a wide range of other scientific purposes e.g. space weather forecasts and detecting space debris.

EISCAT_3D users are expected to access the User Analysis Facility through a user portal (Web) or a command-line interface to the virtualized resources. Metadata searches for analyses may also be performed through either the EISCAT_3D portal or a command line interface. The data to be analysed must be accessed from the data centres from the fast and slow data stores and transferred to the computing resources where the analysis code will run. As the EISCAT_3D users will access the computing e-infrastructure from different countries (also expected to be from outside the Nordic area), a common means of authenticating (identifying) users and authorising access is needed.

EISCAT_3D has a large set of user roles and functions. They have developed tools to share datasets and provide access to computing resources by means of a Master Portal. The Portal currently provides access to selected datasets by means of IP-based Authentication and Authorisation of users.

The goal of the pilot has been to onboard the EISCAT-3D community to embrace Federated AAI and eduGAIN and move away from the IP based Authentication model they are currently using for their portal. The newly built AAI should be based on the AARC BPA for authenticating and authorising their users.

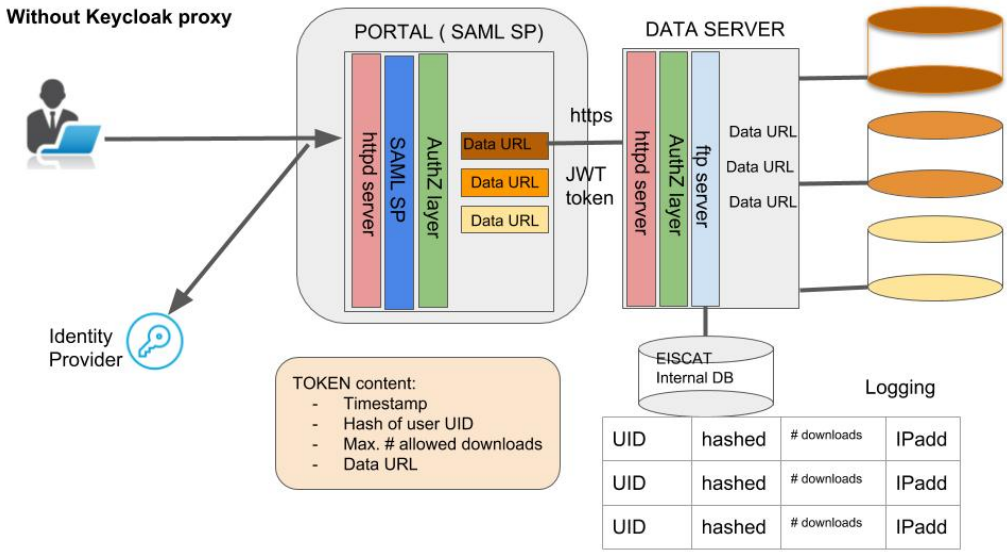

Figure 4: Architecture based on BPA for EISCAT_3D

3.3.1 Technical description

To ensure moving from IP based Authentication to federated AAI, a two tiered system has been design and implemented based on a user-facing Master portal and a data-facing Data Server. Both these tiers were realized in Python language.

In addition to the 2 layers, an internal EISCAT database has been set up to store authentication specific information for individual users. The Database has been used to implement access control to the data sets for users, like enforcing a maximum number of allowed download for a specific. A complete docker compose based deployment suite has been developed, and specific containers have been provided to implement the Master Portal, the DataServer and the Database. Docker containers are built by means of a specific customized Docker build file and subsequently deployed. The Authentication layer is implemented using the lazy sessions model on the portal, where specific bits of the portal are accessible only to authenticated users. Once logged in, a user is presented with a data set selection table, and URLs exposing the data sets included a comprehensive token to uniquely identify the user and allow for logging her/his datasets download activity. The token included a timestamp, a hash of the user ID, a data URL to fetch data sets from and the maximum number of allowed downloads.

3.3.2 Pilot Results

The pilot has achieved some of its original goals, including successfully moving from IP-based access to a federated AAI model, and an automated docker-based suite to deploy the whole pilot AAI is available and will be optimised. The pilot successfully demonstrated access to datasets based on federated authentication and authorisation of users.

The current system consisting of a portal and a download server has been replicated onto a new pilot infrastructure providing an IdP, a Service Provider protecting the Master portal, and a Data server exposing the data. The pilot infrastructure thus provides the same functionality to EISCAT_3D users but making use of their Federated Identity. The pilot has been developing a comprehensive Docker-based installation and configuration suite, in order to automate and facilitate its deployment as much as possible.

A fundamental next step will be the inclusion of the KeyCloak proxy to implement the prototypical community proxy for EISCAT_3D which allows multiple Service Providers to become part of their infrastructure without having the need for the individual Service Providers to deal with the complexity that comes with Federated AAI.

EISCAT_3D will be managing its AAI solution in future based on the AARC pilot, so that federated authentication will be central to the management of their users and services. Through the pilot system, all relevant requirements from EISCAT_3D have been addressed and the subsequent evolution of the system will be capable of serving the community at large, including external users to eduGAIN.

3.3.3 Requirements Table

Category	Requirement	How it was addressed	Achieved?
Federated access	Create an IdP and SP for the EISCAT master portal	IdP and SP created. Attribute release has been provided by the locally deployed Identity Provider hosting users for the purpose of the pilot	✅
On-boarding & support	Proxy test/dev environment - a separate environment where new features could be tested	A test environment for EISCAT has been provided to enable hands-on testing with AAI	✅

Table 3: EISCAT_3D requirements

3.4 European Plate Observing System (EPOS)

EPOS is a pan-European collaboration which aims to establish a comprehensive multidisciplinary research platform for the Earth sciences in Europe. It spans 25 countries, involves four international organisations and 256 national research infrastructures. The expected number of users will likely grow to a total of 2000. EPOS has already established an AAI prototype with the EGI CheckIn service as an IdP and Unity-IdM as its core. During the AARC pilot, the existing infrastructure was vastly extended to meet the full EPOS requirements concerning AAI and establish a more mature, production setup.

The ultimate goal for EPOS is to implement SSO for their users while accessing EPOS services: these include the so-called EPOS Thematic Core Services (TCS) – web-based services in specific Earth Science domains – and Integrated Core Services (ICS) – general, cross-domain computing and storage resources, user management, metadata catalogue. TCS and ICS are interconnected by an EPOS interoperability layer.

The abstraction, interoperability layer will ensure interoperation between the Integrated Core Services and the Thematic Core Services. As an example, Cloud services can be provided at the ICS level, but in some cases need to be accessed by a TCS. National Research Infrastructures will contribute to the provisioning of TCS services.

The challenge for EPOS was mainly dealing with a huge community scattered across different countries and institutions who were already used to offering services in a certain way. The newly built AAI had to support existing users, authorisations and workflows. Moreover, the Integrated Core Services which are offered centrally were still in development during the pilot; thus, the decentralised Thematic Core Services were not used to dealing with a central component.

To accomplish this, EPOS designed an architecture based on the AARC BPA with a centralised proxy that also acts as an integrator between the different Thematic Core Services. Each Thematic Core Service has its own proxy which allows delegation towards other Thematic Core Services. In practice this means that users and attributes stored locally at one Thematic Core Service can, if they have the correct authorisation information, also access another Thematic Core Service. Legacy Thematic Core Services can use EPOS's centralised proxy to accomplish the same result but without the need to implement a complex proxy themselves.

3.4.1 Pilot Results

The goals of the pilot were specified on several levels ranging from networking to technical ones. All goals identified and achieved are listed below:

1. Strengthening collaboration between EPOS and e-infrastructures.
2. Simplify access for EPOS services and EPOS uses to computing and storage cloud services offered by e-infrastructures.
3. Validation of EPOS AAI model and chosen solution with AAI experts, possibly obtaining direction and recommendations for future evolution.

4. Provide technical feasibility study based on implementation of integration between an actual EPOS component and an existing e-infrastructure service.

Such goals serve the mutual interests of both EPOS and AARC technology specialists. The EPOS community has gained a framework of collaboration and technical guidance for the important components of its solution. The AARC community has gained another use case to validate the existing blueprint architecture and gather ideas about the challenges related to AAI.

Goals 1, 2 and 3 described above were fulfilled during the AARC2 project workshops and trainings given by AARC2 for EPOS community members. Dedicated sessions and small group virtual meetings were held to discuss the details of EPOS's approach to AAI.

Technical implementation related to the 4th goal was focused on the case described in the next section.

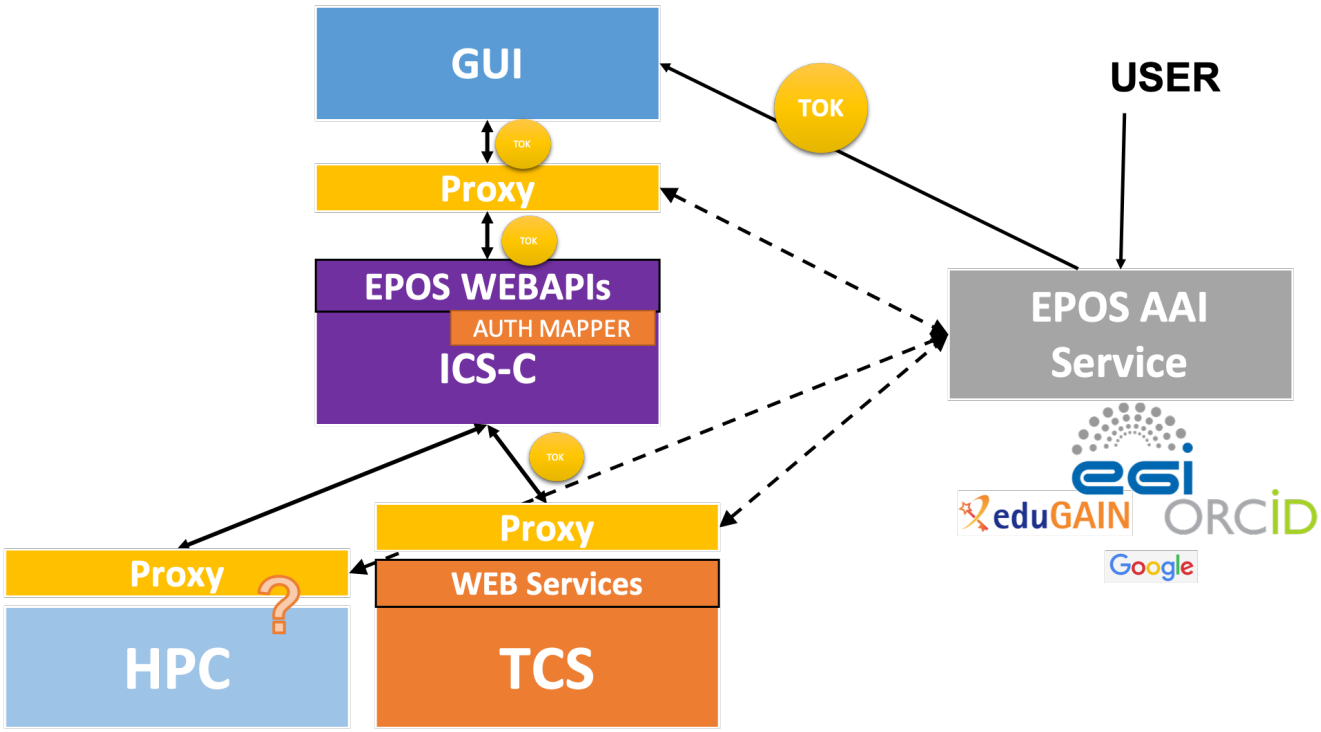

Figure 5: Architecture based on BPA for EPOS

3.4.2 Requirements Table

Category	Requirement	How it was addressed	Achieved?
Identity lifecycle & linking	Account linking	The ability, for one entity, to link credentials from multiple IdPs to one account on an SP. More generically, the ability for a researcher to link multiple identities together, whether held in	✅ .

		parallel or succession. The ability to accurately link accounts depends strongly upon the release of an appropriately unique and persistent identifier.	
AuthZ	Real-time authorisation AuthZ decisions at an SP must be based on identity credentials, attributes or assertions that have a short lifetime, i.e. they are valid now and not for too long into the future. Even within this short period it should be possible for the SP to look up real time status information, e.g. revocation lists and/or suspension lists.		✅
	Group management	Research communities must be able to add individuals to Groups, for use in AuthZ, Quota management and Accounting. Groups should be hierarchical, and users can belong to more than one group.	✅

Table 4: EPOS requirements

3.5 LIGO

The LIGO Scientific Collaboration (LSC) is a group of scientists focused on the direct detection of gravitational waves, using them to explore the fundamental physics of gravity, and developing the emerging field of gravitational wave science as a tool of astronomical discovery. The LSC works toward this goal through research on, and development of techniques for, gravitational wave detection; and the development, commissioning and exploitation of gravitational wave detectors. The LSC carries out the science of the LIGO Observatories, located in Hanford, Washington and Livingston, Louisiana as well as that of the GEO600 detector in Hannover, Germany. Our collaboration is organised around three general areas of research: analysis of LIGO and GEO data searching for gravitational waves from astrophysical sources, detector operations and characterisation, and development of future large-scale gravitational wave detectors. Founded in 1997, the LSC is currently made up of more than 1200 scientists from over 108 institutions and 18 countries worldwide.

Each member of the LSC is assigned an albert.einstein identity and they manage this account and their credentials via the my.ligo.org application. This pilot aims to investigate the infrastructure and organisational changes required to support the use of federated institutional entities alongside existing internal credentials. In particular it will identify technological components and deploy a pilot service to be used for evaluation. It will also work to understand the current limitations of federated identities as applied to the LSC and recommend alternative approaches where applicable.

SAML proxies are increasingly being used to easily connect all of a collaboration's resources into the eduGAIN network and this would demonstrate it's success for a large, established collaboration. The AARC Blueprint Architecture is important in shaping the design and features of this pilot.

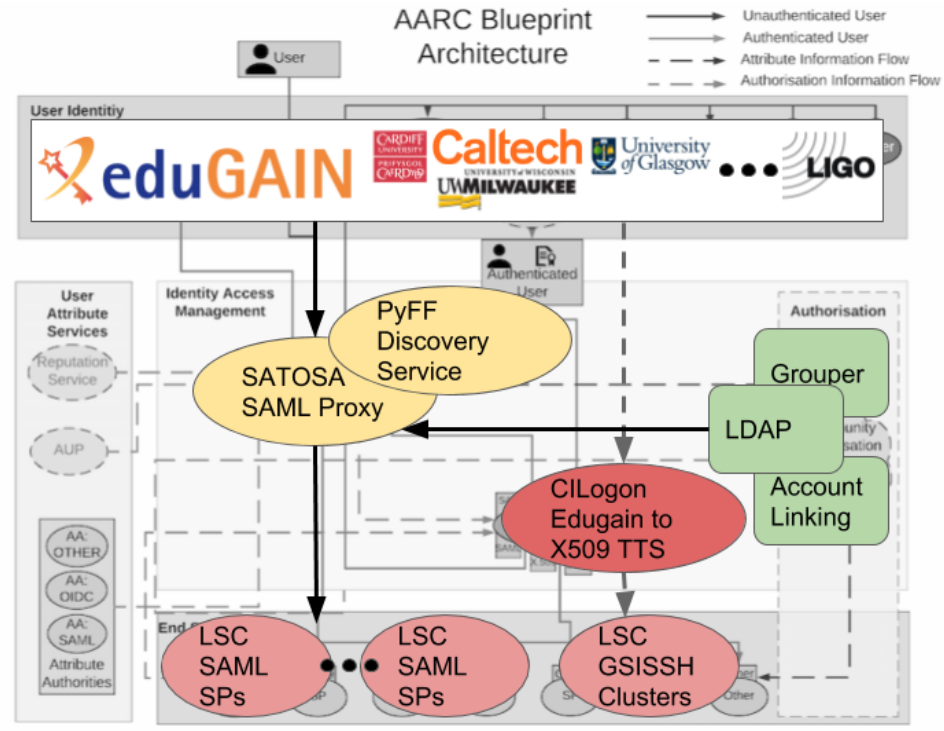

Figure 6: LIGO architecture based on the AARC BPA

3.5.1 Technical description

The aim of the pilot is to investigate the use of a SAML proxy to enable adoption of federated identities. As part of the AARC2 pilot we have deployed an IdP/SP proxy using SATOSA software and COmanage was used to manage the groups.

The proxy frontend (IdP) is registered with LIGO Metadata and the backend (SP) is registered with eduGAIN. The proxy uses existing microservices to process authentication response in the following way:

- First pick up correct identity attribute from differing IdPs
- Use identity attribute to lookup information in LIGO LDAP for example
    - uid: paul.hopkins
    - sn: Hopkins
    - givenName: Paul
    - mail: paul.hopkins@ligo.org
    - isMemberOf: ….
    - …
- Correctly reproduces attribute behaviour of current system

For the pilot PyFF Discovery services (it is an implementation based on the RA21 guidelines) is used for metadata aggregation and login chooser.

3.5.2 Pilot Results

A pilot instance was deployed and registered in the eduGAIN metadata and underwent extensive testing using a number of existing LSC resources. Within the pilot, account linking between institutional identities and a user LSC identity was performed using a manual administration step.

Going forward, more work is planned to move the pilot into production; this work includes:

- Register the SAML Proxy SP as R&S; which will allow IdPs to release more attributes
- Complete configuration of PyFF for metadata aggregation and discovery service
- Enable account linking service
- Research replacement for SAML ECP
- Work closely with LIGO Identity and Access Management group to deploy the proxy in a fault tolerant manner.

3.5.3 Requirements Table

Category	Requirement	Explanation	Achieved?
Identity lifecycle & linking	Account linking The ability, for one entity, to link credentials from multiple IdPs to one account on an SP. More generically, the ability for a researcher to link multiple identities together, whether held in parallel or succession. The ability to accurately link accounts depends strongly upon the release of an appropriately unique and persistent identifier.	LIGO users able to link their internal account to their institutional account, using existing interfaces. New tools and services will be required to enable account linking for new accounts.	✓

Discovery & usability	Smart discovery IdP discovery should be “smart enough” to quickly and easily take a user to their appropriate home IdP. For example show the user a short list tailored to them by home country, institute, e-Infrastructure, research community, project or other hints.	The Discovery Service can restrict the list of displayed IdPs to those that host LIGO members. This is currently around 110. The Discovery Service can also use other hints to aid discovery at observatory locations.	
	Logo in metadata at an agreed standard size Discovery services should display organisation logos to aid the user in choosing the IdP. IdPs or research community proxies should provide a logo at an agreed standard size.	To aid users in choosing their home institution the icon is displayed in the list. Where icons don't exist or don't conform to an acceptable size then their can be manually changed.	
Attribute release	Attribute release IdPs must release a unique, persistent, omnidirectional identifier, email address, and name for users when accessing research services. For example, ensure that the CoCo and R&S entity categories are widely adopted.	By adopting R&S the proxy is able to make use of a email address for linking accounts which can simplify account linking using the existing tools and services.	
	Attribute release across borders The R&S bundle, especially, needs to	By adopting R&S the proxy is able to make use of a email address for linking accounts which can simplify account linking using the existing tools and services.	

	easily flow from IdPs to SPs without regard to their nationalities. More outreach of the risk analyses performed by GEANT and REFEDS about R&S + CoCo entity categories is needed to increase adoption.		
Research community proxies	IdP/SP Proxies must be allowed to join edugain IdP/SP Proxies must be permitted to join eduGAIN or one of its constituent national federations. Snctfi requirement below is related.	The IdP/SP proxy was added to eduGAIN and assigned R&S status.	
Interoperability	Avoid user/interop issues due to inconsistent propagation of metadata for entities. Federations should support standard and automated metadata propagation processes and, where out of band actions are required, provide clear documentation and support	The attributes required to add join eduGAIN and R&S status were found using the Internet2 documentation.	✅
	Federation entity attributes designed to enhance user experience should be populated Eg, the entity attributes defined in	Security and error information was added to the eduGAIN metadata for the proxy.	✅

	the SAML "MDUI Information" specification and errorURL should be populated at least.		

Table 5: LIGO Requirements

3.6 Life Science

The Life Science (LS) community includes 13 different research collaborations and e-infrastructures that clustered in the CORBEL initiative (Coordinated Research Infrastructures Building Enduring Life-science Services) that is charted to create a platform for harmonised user access to biological and medical technologies, biological samples and data services required by cutting-edge biomedical research in a coordinated manner. One of the goals of the community was to operate one single AAI for the benefit of the whole Life Science disciplines.

With this aim in mind, a sub-set of these research infrastructures, namely EMBL, BBMRI-ERIC, EMBL , Infrafrontier and Instruct) joined the AARC2 project to pilot an AAI that would follow that AARC BPA.

This community had very clear requirements [LSrequirements] and some of them had already a clear understanding of what operating an AAI entails. To ensure long term sustainability and given that running an AAI is not a core business for any research collaboration, they agreed to outsource the operations to the e-Infrastructures (EGI, EUDAT and GÉANT). A call for proposal based on the LS requirements was issued and in April 2017 EGI, EUDAT and GEANT submitted a joint proposal. The pilot phase was run as part of the AARC project, as effort for a similar pilot was already allocated.

3.6.1 Technical description

The joint e-infrastructure proposal was based on the architecture that is depicted in the Fig 6 below.

The architecture follows that AARC BPA model; characteristics of the pilots were:

- the implementation the operation and functionalities of the proxy are split across three operators
  - GÉANT operated the interface to the IdPs to enable federated authentication via eduGAIN;
  - EGI operated he interface to the SAML services
  - and EUDAT operated the interface to the OIDC services;
- the LS community chose to use PERUN as their group management system, as many of them were already using it. Upon federated authentication of the user, users are assigned to a group via Perun, which offers support for full life cycle of user identities. This includes the initial user registration, the acceptance of the terms of use, account linking, support for user consent for the release of user attributes, group and role management, delegation of administration of groups to authorised users and the configuration of custom enrolment flows for groups via an intuitive web interface.
- the community required a hostel IdP for users without a federated account; a similar service was already operated by some of the LS communities;

- WaTTS (the INDIGO Token Translation Service) was used to provide token translation functionalities. WaTTS is a plugin-based Token Translation Service and translation to various credentials (e.g. SSH keys, X.509, S3 access tokens, etc.) are supported.

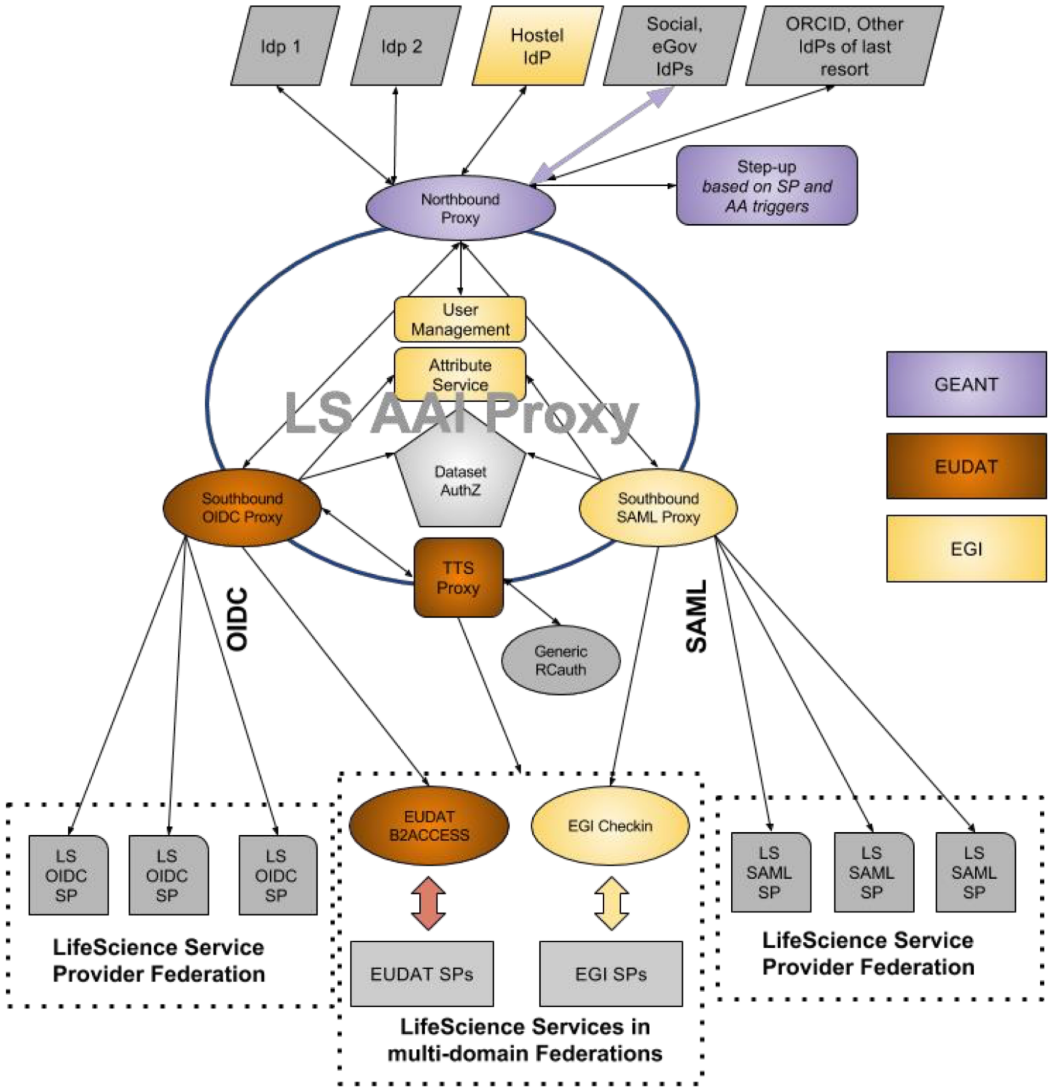

Figure 7: LS AAI Architecture based on the AARC BPA

3.6.2 Pilot Results

The LS AAI started the pilot in AARC2 with the aim to use the results to deploy a production infrastructure; provision for the deployment was secured via a different EC-funded project, EOSC Life (started in March 2018).

The pilot in AARC2 went beyond the technical aspects and also explored operational aspects related to outsourcing the operation of an AAI to multiple e-infrastructures, availability, monitoring and so on. This offered

also an opportunity for e-infrastructure to better understand how they can collaborate and how to offer an AAI as a service. The AARC BPA proved to be ARC BPA flexible enough to cater for complex use-cases and to be deployed in a multi-operator model. Although they final architecture for the production LS AAI may be different, AARC2 pilot results have been useful for the LS community to understand how to move forward.

3.6.3 Pilot Results

The table below shows the main technical requirements that were implemented in the LS AAI AARC2 pilot.

Category	Requirement	How it was addressed	Achieved?
Identity and identifiers	Life Science ID	A persistent, non-re-assignable identifier needed targeted to the Life Science community. This is generated in Perun using identity information provided by the proxy upon authentication	✓
	User Identifiers	These are the lifeScience ID and the lifescience user name, which are stored in Perun.	✓
Registering and authenticating	Supported Authentication Providers	Support all identities coming via eduGAIN, via Research infrastructures AAIs, Commercial providers (such as, Google), ORCID and Hostel Identity Provider	✓
	Hostel Identity Provider	Support users who cannot use any other Authentication providers listed in the previous section.	✓
	Account linking	A user can link multiple accounts from multiple authentication providers (see the previous section) to their Life Science user ID, at first logging in using a previously linked account and subsequently the new account, or by demonstrating control of an e-mail account, using a procedure that is as secure as above.	✓
	Account management	ach Life Science service ID must have at least one associated Life Science user ID that belongs to a natural person who manages the account and takes responsibility of the activity done using the service ID.	✓

	Assurance Framework	Develop a framework based on REFEDS Assurance Framework to manage the assurance for: 1. identity proofing (registration, Identity vetting), 2. authentication (login) and 3. attributes	
	Step-up Authentication	Allow a user to associate a second authentication factor to their Life Science ID and a Relying service can ask the Life Science AAI to perform a step-up authentication using it.	
Attributes and authorisation	Home Organisation Affiliation(s) of a user	Each user can be affiliated to one or more Home organisations (such as, a university, research institution or private company) and the user's affiliations may change over time. A Relying service wanting to couple user's permissions to their continuing affiliation can observe the Home Organisation Affiliation attribute and their changes.	✅
	User's Research Infrastructure attribute	User's Research Infrastructures attribute indicates to which research infrastructures the user's Home Organisation is affiliated with.	✅
	Groups	The Life Science AAI requires a service for managing users' group memberships and roles in the groups they belong to. Management of groups is done using a web interface. Each user can belong to one or several groups.	✅
	Other attributes	The Life Science AAI should allow for arbitrary attributes to a Life Science ID, including: ● If the Life Science ID is a user ID or a service ID ● User's name ● User's e-mail address (which is confirmed by an e-mail handshake) ● User's ORCID ID (which is recorded using ORCID APIs) Other wider researcher identifiers (such as, a researcher ID assigned to users by e-infrastructures) if they emerge	✅

Interfaces	Credential translation	Relying Services can subscribe to the credential translation service of the Life Science AAI, allowing the users to obtain X.509 certificates based on the login described in the previous section	✅
General statistics	The LS AAI should provide anonymised statistics on #of Relying services, #of identities, # of logins	The proposed pilot has all elements in place to provide the necessary statistics.	✅
Security incident response	The proxy must commit to Sirtfi.	Work started but it was not concluded in the AARC pilot.	

Table 6: LS AAI requirements

3.7 LifeWatch

LifeWatch-ERIC is a European Research Infrastructure Consortium providing e-Science research facilities to scientists seeking to increase their knowledge and deepen their understanding of Biodiversity organisation and Ecosystem functions and services in order to support civil society in addressing key planetary challenges. LifeWatch-ERIC was established as a European Research Infrastructure Consortium by the European Commission Implementing Decision (EU) 2017/499 of 17 March 2017. During its ESFRI stage, LifeWatch was composed by different national initiatives working on different services and solutions for the research community. During this new ERIC stage, LifeWatch ERIC requires a solution to provide access to the different services in a common way, as well as organise the different groups and roles defined. Currently, the different LifeWatch services, Virtual Laboratories and Virtual Research Environment manage their own local users, with some exceptions that allows institutional IDs. The underlying technologies depends on the services, but they mainly support web-based authentication, with some exceptions, for example using HPC resources.

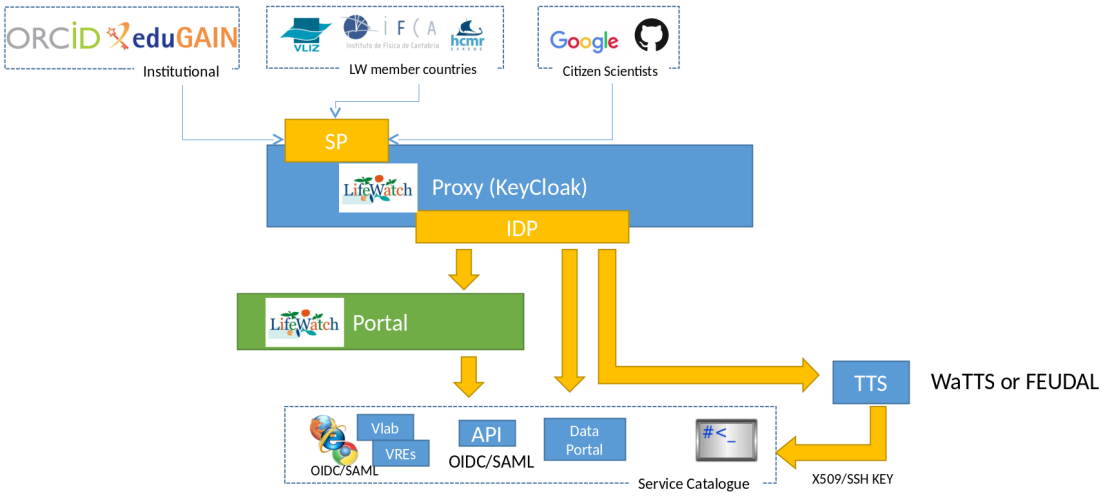

Figure 8: Architecture based on the AARC BPA for LifeWatch

This pilot activity aims to identify and enhance an existing AAI solution to be adopted by LifeWatch ERIC as IdP, integrating already existing institutional or social identities in a federated way.

This IdP solution will be used for the following purposes:

- To give access to restricted LifeWatch services. The services may be restricted because of processing power or storage demands.
- To protect user data and scripts that are stored on the infrastructure (Unix home folders etc.).
- To give access to data not yet in the public domain (data in databases, project moratorium period).
- To distinguish between users uploading data to the system (RvLab , eLab, data explorer).
- To give access to OpenStack configuration interface and computing resources at the infrastructure layer.
- To manage roles/groups and authorise them to access specific services.
- Currently, the different user apps manage their own users. The institutional credentials could be federated with the users' Identity Provider. It should also manage the following roles/users:
    - IT administrators who have access at infrastructural level;
    - Developers/Solvers who have access to computing/storage resources to develop new Vlabs/VREs;
    - LifeWatch ERIC research users;
    - Citizen Science (to have access to applications).

The architecture suggested by AARC based on its blueprint is a promising approach to be adapted to the European framework for the European Open Science Cloud.

3.7.1 Technical description

During the evaluation phase, different components were checked, including EGI check-in, B2ACCESS, INDIGO IAM, and KeyCloak, to decide which suited the LifeWatch ERIC needs better. Finally, KeyCloak, an open source

solution supported by RedHat and adopted by different communities, was chosen. The reason for selecting KeyCloak was the set of features it provides, which are sufficient for the needs of LifeWatch ERIC as a community, including:

- Different user management options. KeyCloak [KeyCloak] allows to create a local user in a database or connect to a LDAP systems.
- User federation using the main technologies: OpenID Connect, SAML, Oauth2. It is pre-configured with many different social IDs (Google, Facebook, GitHub, other KeyCloak instances) and eduGAIN.
- Unlimited federation of IdPs which is needed for the complex LifeWatch community, with representatives from many institutions.
- Customisable set of attributes, both for local users as well as federated ones.
- Attributes mapping from federated IdPs. This is needed to classify the different roles expected.
- Group and role management to identify user permissions.
- Easy to install and maintain. It works with a database that can be distributed.
- Clustered mode to set up a high-availability environment.

3.7.2 Pilot Results

The pilot has been implemented and deployed in a testbed to demonstrate that everything will work as expected. The AARC BPA has been used to identify which components are needed to address the pilot needs.

The pilot will be the official LifeWatch ERIC IdP and it will be used to access the services, taking into account the different roles in the community. It will be deployed in a high-availability environment since it will be a critical service for the Research Infrastructure, and it will be one of the keys to integrating LifeWatch ERIC in the context of the European Open Science Cloud, so that the sustainability of the pilot is guaranteed.

3.7.3 Requirements Table

Category	Requirement	Explanation	Achieved?
Identity lifecycle & linking	Account linking - The ability to link multiple identities together, whether held in parallel or succession. It depends strongly upon the release of an appropriately unique and persistent identifier.	Keycloak allows the account linking for users accessing with different Identity providers. Once the user is logged in with any of the IdPs, any other can be selected to login again in order to link both.	

	ORCID - provide login at ORCID IdP	ORCID includes Open ID Connect clients that can be added as a federated Identity Provider in Keycloak. It has been succesfuly tested.	✅
AuthZ	User blocking - the ability to block access to a service	The users can be totally (account disabling) or partially (based on attributes or roles) blocked thanks to the keycloak administration panel.	✅
	Deprovisioning - the ability to suspend or remove an individual's access when they no longer possess right of access	The users can be totally (account disabling) or partially (based on attributes or roles) deprovisioned thanks to the keycloak administration panel.	✅
	Group management - the rights to add individuals to hierarchical groups	Keycloak has the capacity of creating roles with specific rights as well as different types of groups and subgroups.	✅
	Active role selection - select which attributes, groups or role are active for a specific connection request	Based on attributes provided by the different federated IdPs, keycloak assign the logged in user to specific roles or groups.	✅
Security incident response	IdP suspension - the Ability to disable all logins from identified IdPs as part of managing a security incident. Can happen by home federation or by Proxy.	If there is any security issue, one IdP can be easily disconnected (temporal or not) from keycloak.	✅
On-boarding & support	eduGAIN test/dev environment - Easy-to-use testing environments to allow new Proxies and new SPs to experiment with their Federation facing parts without interfering with existing	Keycloak offers the possibility to provide different environments called realms, where different configurations (development or production) can coexist.	✅

	production deployments.		
	Proxy test/dev environment - a separate environment where new features could be tested separately from production	Keycloak offers the possibility to provide different environments called realms, where different configurations (development or production) can coexist.	
Sustaining critical infrastructure	IdP of Last Resort - Provide sustained services to meet the many cases where global researchers do not have access to an acceptable Home Organization IdP	Apart from IdP federation, Keycloak can manage users internally or linked to other types of user management systems like LDAP.	

Table 7: LifeWatch requirements

3.8 Worldwide LHC Computing Grid (WLCG)

WLCG has been operating a distributed computing infrastructure for the past 15 years. User authentication and group management is based on X.509 certificates, with authorisation conveyed in VOMS Proxy certificates. This is no longer considered best practice, both in terms of user experience and of infrastructure sustainability since the community at large is moving to OAuth 2.0 token-based authentication and authorization models.

This pilot activity aims to identify and enhance existing AAI services to suit the requirements of the High Energy Physics community. The requirements focus on aspects currently not included in AAIs, a sample of which are listed here:

- A user-friendly workflow to provision authorisation tokens in the user's local environment for command line activities. The majority of physicists' time is spent submitting "jobs" (analysis code) from a terminal and it is essential that limited browser interaction is required for authentication/authorisation.
- Integration with existing infrastructure for a smooth transition. Token translation to and from X.509 certificates will be essential for backwards compatibility. The existing database of identity vetting information should also be leveraged.
- Development of a shared JWT profile for the wider physics community (and/or beyond)

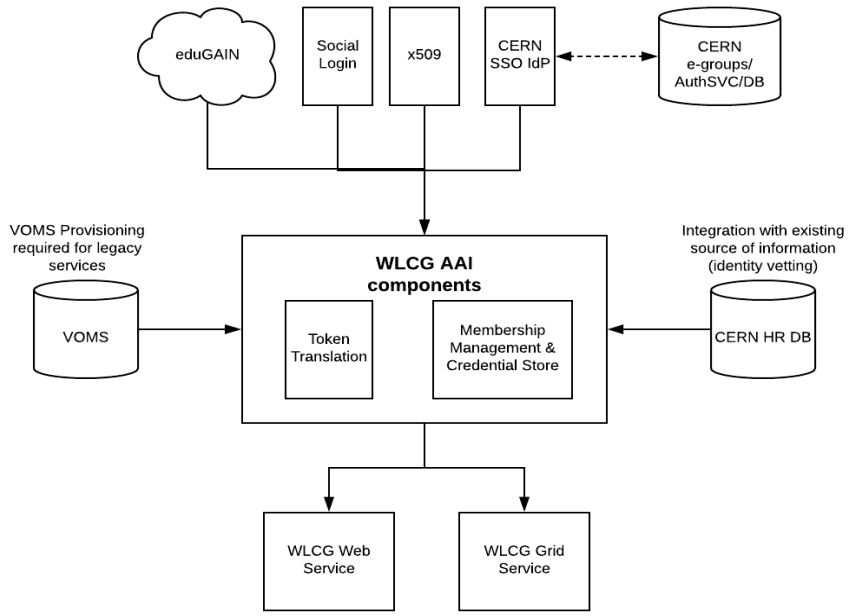

Figure 9: Architecture based on BPA for WLCG

A priority for WLCG was to not reinvent the wheel, following the FIM4R recommendation to re-use shared components. EGI Check-in was identified by the AARC project as an existing solution that already fulfilled many WLCG requirements, and was chosen for enhancement during this pilot. The RCAuth.eu service was highlighted as another critical component to enable workflows required by the High Energy Physics Community.

3.8.1 Technical description

A WLCG instance of the EGI Check-in codebase (SimpleSAMLPhp, COManage and MitreId) was deployed on multiple Virtual Machines on CERN's openstack infrastructure behind an nginx proxy. CERN's internal User Identity Vetting source (the Human Resources Database) was integrated through a REST API layer developed by collaborators from INDIGO-IAM. Several VOMS plugins were developed for COManage, to facilitate backwards compatibility. The WLCG Check-in instance was registered as a relying party for RCAuth.eu to enable provisioning of IOTA certificates. A ssh client was deployed to provide token provisioning in the user's local environment.

3.8.2 Pilot results

An enhanced version of the EGI Check-in code was successfully deployed by AARC team members at CERN and able to satisfactorily manage WLCG membership workflows and token provisioning. Certain aspects of the pilot (particularly token translation) will be included in a production AAI over the coming years, the roadmap of which remains to be clarified. AARC participation was critical in the development of a JWT Schema for WLCG Tokens.

A key aspect to note is the interest shown by other Physics communities who are also looking to use existing AAIs for their experiments. The pilot was directly useful in providing prototypes, proofs of concept, and demonstrations.

3.8.3 Requirements Table

The following table includes requirements for enhancing EGI Check-in. Any requirements that were already satisfied are not included.

Category	Requirement	Explanation	Achieved?
Identity lifecycle & linking	Account linking The ability, for one entity, to link credentials from multiple IdPs to one account on an SP. More generically, the ability for a researcher to link multiple identities together, whether held in parallel or succession. The ability to accurately link accounts depends strongly upon the release of an appropriately unique and persistent identifier.	Users are able to link SAML to X.509, with X.509 certificates generated by RAuth.eu	✅
AuthZ	Group management - Research communities must be able to add individuals to Groups, for use in AuthZ, Quota management and Accounting. Groups should be hierarchical and users can belong to more than one group.	WLCG requirements for user and group management were configured in COManage.	✅
	Active role selection - Individual users must be able to select which attributes, groups or roles are “active” for a particular connection request and AuthZ decision.	Users are able to specify whether their session is active in the scope of a particular role, for example as either a researcher or as a system administrator.	✅
	Integration of Experiment Identity Vetting - Identity Vetting performed by Experiment Secretariats should be integrated into the AAI to contribute to combined assurance	Users’ membership requests are checked against the existing source of Identity Vetting Information (HR Database). The experiment affiliation and dates are extracted and used to validate the user’s account in the system.	✅

Non web	Non-web use cases & support - A very important requirement for research communities. Many interactions between clients and servers are via the user's command-line or via interacting applications using API access to AAI. Cannot assume that all access will be via a web browser interface, or that a web browser will be part of the authentication flow, even beforehand to set things up. Strong authentication (not necessarily MFA) may be required for some use cases.	Researchers are able to upload an ssh key into COManage and run a script locally to provision X.509 and OIDC tokens in their local environment.	✅
	Credential translation - Services will not always be able to consume the credentials the user currently has. Translations from one type of credential to another is a very common and important requirement.	RCAuth.eu has been integrated to allow token translation to X.509 based on user attributes collected in COManage.	✅

Table 8: WLCG requirements

4 Conclusions

As the document described, AARC2 worked closely with research infrastructures interested in testing the deployment of an BPA compliant AAI. The results of the various pilots in SA1 demonstrated that the AARC BPA works in practice and can integrate with existing components that each community may already in production. The AARC BPA being a blueprint is not prescriptive about technologies and tools that communities may wish to use, although the team was able to provide more hands-on support on some tools rather than on others. However, it was entirely to the descrition of the research communities to decide what tools to use to implement the various elements of the BPA.

SA1 also demonstrated the steps needed for implementing an AAI. The requirements collection is a critical aspects and it may take a considerable amount of time. The life science community have set an example, and their requirements have been used as starting point by other communities as well. Once the requirements are known, the design phase can start.

In most of the AARC2 pilots, after this initial "requirements analysis" phase, the selected components were installed and configured in a test setup. Once this phase had been completed and whole AAI environment was

ready, tests could be carried out with production services and with a selected number of users to verify whether the architecture worked in practice and whether the initial requirements were covered ("evaluation"). In the last phase ("finalisation"), all documentation was made available.

The transition to production was out of scope of the AARC2 pilot; this was in line with what described in the AARC2 description of work. The reason for this being that typically each research infrastructure has their own internal processes and timelines to migrate new software into production. With the exception of the life science community that from the beginning of AARC2 had agreed to ask the e-infrastructures (EGI, EUDAT and GÉANT) to operate their AAI on their behalf, all the other communities indicated their intention to operate their own AAI.

The learning path to migrate to an AARC BPA compliant AAI is clearly shorter for those communities that have already an operational AAI in place and that are already supporting federated access.

Important success factors for implementing an AAI are the involvement of key Service Provider operators from the start and to work with well-defined requirements, as well as conducting careful and constructive discussions within the user communities. Moreover, providing support for policy aspects and trainings has proved to be essential in achieving the expected results.

References

[AARC_BPA]	https://aarc-project.eu/architecture/
[COmanage]	https://www.internet2.edu/products-services/trust-identity/comanage/
[Grouper}	https://www.internet2.edu/products-services/trust-identity/grouper/
[KeyCloack]	https://www.keycloak.org/
[LSrequirements]	https://goo.gl/zvTQmB
[FIM4Rv2]	http://doi.org/10.5281/zenodo.1296031
[Satosa]	https://github.com/IdentityPython/SATOSA
[Sirtfi]	https://refeds.org/sirtfi
[SNCTFI]	https://www.igtf.net/snctfi/igtf-snctfi-1.0-20170723.pdf

Glossary

AAI	Authentication and Authorisation Infrastructure
AAI service	A service that enables authenticated and authorised access to resources
AUP	Acceptable Use Policy
CA	Certification Authority
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 User Attribute 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 [SIRTFI]).

Community service	A service provided only to members of a specific community
Credential	A set of data presented as evidence of a claimed identity and/or entitlements [X.1254]
Credential Service Provider	A trusted actor that issues and/or manages credentials [X.1254]. In the context of the [RAF] specification, this term refers to the Identity Provider and the associated Identity Management system that manages the user identities and attributes observed by the Relying Parties.
CSP	Credential Service Provider
Digital identity	Information that represents an entity (subject) within a domain. It contains information about the subject's attributes and relationships
ECP	Enhanced Client or Proxy
eduGAIN	International interfederation service interconnecting research and education
EPOS	European Plate Observing System
ERIC	European Research Infrastructure Consortium
ESFRI	European Strategy Forum on Research Infrastructures
Generic service	A service provided to members of different communities
GDPR	General Data Protection Regulation
ICS	Integrated Core Services in EPOS
IdP	Identity Provider
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 membership management. Specifically, the infrastructure proxy comprises two (2) AARC BPA component layers: the Access Protocol Translation and the Authorisation.
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
KeyCloack	Open Source Identity and Access Management Solution
OIDC	OpenID Connect
R&S	Research and scholarship
Relying Party	Actor that relies on an identity assertion or claim
SAML	Security Assertion Markup Language
SATOSA	Open Source proxy for translating between different authentication protocols such as SAML2, OpenID Connect and OAuth2
SP	Service Provider
SSO	Single Sign On
TCS	Thematic Core Services in EPOS
VO	Virtual Organisation
WLCG	Worldwide LHC Computing Grid (WLCG)