# AARC Blueprint Architecture

| | |
|---|---|
| Published Date: | 18-04-2017 |
| Revision: | 1.0 |
| | |
| Work Package: | JRA1 |
| Document Code: | AARC-BPA-2017 |
| Document URL: | https://aarc-project.eu/blueprint-architecture |

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

# 1 Introduction

The way researchers collaborate within scientific communities can vary significantly from community to community. On the one hand, there are highly structured communities with thousands of researchers who can be virtually anywhere in the world. These communities have been working together for a long time, they want to share and have access to a wide range of resources, and they have had to put in place practical solutions to make the collaborations work. On the other hand, one finds a number of smaller and more diverse research communities working within specific or across scientific disciplines. Typically, these are either nascent communities being established around new scientific domains, or they are communities in specific domains that did not have to promote widespread and close collaboration among the researchers. And of course, in between these two extremes, there are many scientific communities of varying size, structure, history etc.

During the last two years of the AARC project we have been working together with e-infrastructures, research infrastructures, research communities, AAI architects, and implementers to get a better understanding of their experiences and needs in sharing and accessing resources within research collaborations. Our goal has been to collectively define a set of architectural building blocks and implementation patterns, which we call the 'AARC Blueprint Architecture', that will allow the development of interoperable technical solutions for international intra- and inter-disciplinary research collaborations.

The first version of the AARC Blueprint Architecture [AARC-BPA-2016] was published in July 2016. In that document, we analysed the architectures of existing designs and implementations and extracted a high-level architecture that encapsulated common patterns and building blocks. The second version [AARC-BPA-2017] was published in April 2017 and builds upon AARC-BPA-2016. AARC-BPA-2017 extends the previous version and provides guidance on topics such as access to non-web based services, token translation services (TTS), best practices for managing authorization, harmonized expression of group membership and role information, attribute aggregation and credential delegation.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

# 2 AARC-BPA-2017

This version of the AARC Blueprint Architecture (AARC-BPA-2017), builds upon the previous one and provides a more detailed layered architecture, while retaining full backwards compatibility. AARC-BPA-2017 retains the same four layers, each of which includes one or more functional components, grouped by their complementary functional roles. The User Identities layer and the End Services layer are still there, while the Attribute Enrichment layer has been renamed to User Attributes layer and the Translation layer has been renamed to Identity Access Management (IAM) Layer and has a prominent role in the architecture. In AARC-BPA-2017, we introduce a new layer for the centralised Authorisation.

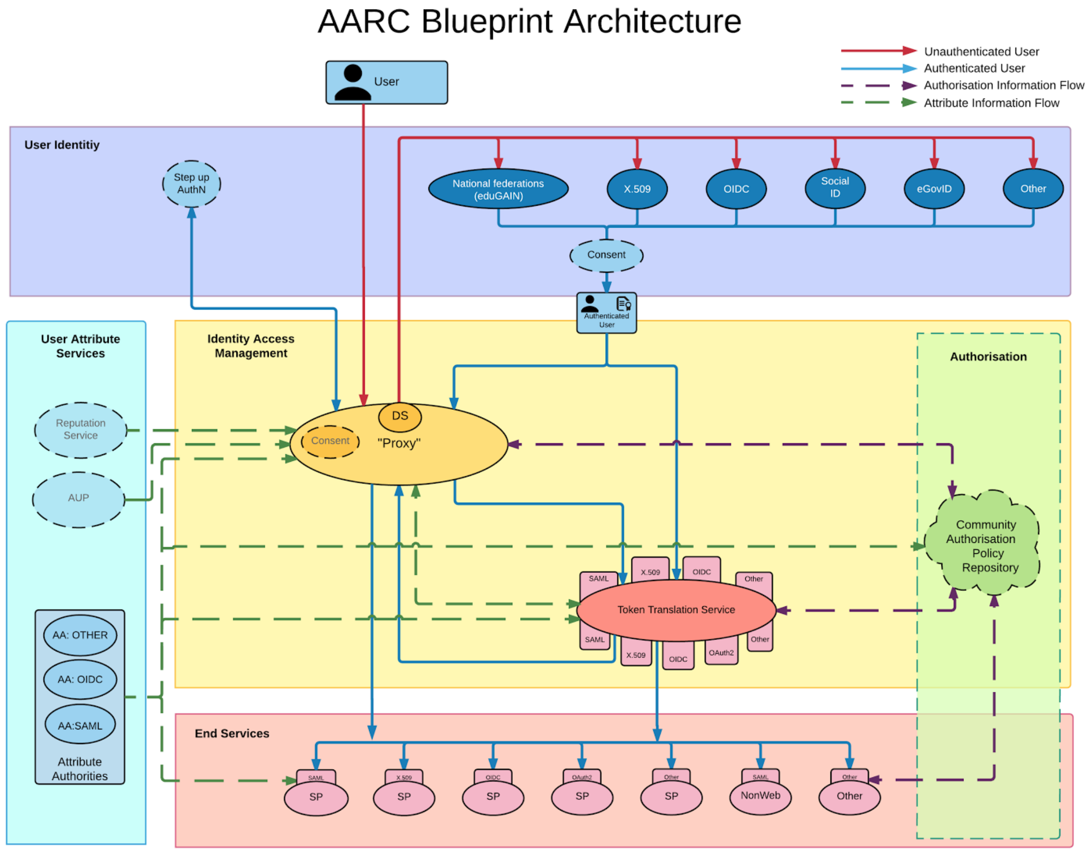

Figure 5: AARC Blueprint Architecture 2017

The **User Identity Layer** includes services which provide electronic identities that can be used by users participating in International Research Collaborations. Typically, identity services in this layer are outside of the administrative boundaries of the International Research Collaborations. Services in the layer are expected to use secure authentication mechanisms, in order to bind physical persons to their electronic identities, which are then made available to services in the other layers via secure protocols.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

The AARC Blueprint Architecture is meant to be technology agnostic by design and has been verified against production implementations that use SAML2, OpenID Connect, OAuth2, or combinations of these. Identity service providers, might utilise single factor or multi factor authentication schemes.

In AARC-BPA-2016, we have incorporated also the traditional flows, which linked directly service and identity providers, in order to denote the different possibilities and the change towards the proxy implementations. In AARC-BPA-2017, we have kept only the flows relevant for the proxy architecture and thus users are shown to access services only via the EI/RI proxy component. In this version of the architecture we also introduced the points where user consent is expected.

The **Identity Access Management Layer** is a new layer, introduced in AARC-BPA-2017, which replaces the Translation Layer. The components in the Identity Access Management Layer are operated by (or on behalf of) the RI/EI and thus reside within the administrative and policy boundaries of the EI/RI. The Identity Access Management Layer, defines an administrative, policy and technical boundary between the internal services and resources of the RI/EI and any other external services and resources. In AARC-BPA-2016, this layer was marked as an optional layer, but in AARC-BPA-2017, the Identity Access Management Layer is an integral part of the architecture as its functionality is required for all the use cases that go beyond the basic Web SSO scenarios.

The components in this layer allow the RIs/EIs to (a) take full advantage of eduGAIN and the national identity federations, (b) reduce the administrative overhead of introducing new services and (c) have the flexibility to choose the appropriate security protocols and mechanisms for delivering internal services to their users. Furthermore, the Identity Access Management Layer enables the implementation of a single point where the RI/EI can provide an IdP discovery service for all its internal services, along with other required functionalities, such as integration with community based group management systems and consistent and scalable central management of user consent. Finally, it enables the infrastructure to provide guarantees that may not be met by the external IdPs alone, such as unique, persistent identifiers, multi-LoA management, infrastructure-specific attributes, account linking, etc.

The Attribute Enrichment Layer has been renamed to **User Attribute Services Layer** and it groups components related to managing and providing information (attributes) about users, such as group memberships and community roles, on top of the information that might be provided directly by the Identity Providers from the User Identities Layer. In AARC-BPA-2017, we have moved this layer on the left side of the Identity Access Management and the End Service layers to make clearer the interactions between components of this layer and the layers of Identity Access Management, Authorisation and the End Services. In contrast with the proxy-managed attributes, responsibility for managing attributes provided by these services rests with the user communities.

Note also the presence of two new components with dotted border, the AUP (Acceptable Use Policy) and Reputation attribute services. The former denotes a service which specifically **records** whether the user has accepted the AUP of the infrastructure (as required by many NRENs, RIs, EIs, and by SIRTFI). The latter is a service which records the user's "reputation" as discussed in [AARC-MJRA1.2]. These components will be further analysed in the next versions of the AARC-BPA in AARC2.

The **Authorisation Layer** is a new layer introduced in AARC-BPA-2017. Authorisation of access to services can take place in many different ways. Typically, in RI/EIs, authorisation can be based (a) on the group membership of the users, (b) on the roles a user might have been granted within the collaboration, (c) on the

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

entitlements users might have been granted, (d) on the affiliations of the users, (e) on the strength of the authentication method used or the quality of the user information or (f) on combinations of these.

Although authorisation enforcement always happens on the service side, the AARC-BPA allows the implementers to delegate much of the complex decisions to central components, which can significantly reduce the complexity of managing authorisation policies, and their evaluation on each service individually. For example, the decision whether a user is allowed to access a specific service can be taken centrally and then communicated to the service by adding a service specific attribute to the user's attributes. In this way, a service can rely on the infrastructure to make the authorisation decision based on a number of appropriate factors. Authorization is a topic which will be thoroughly analysed in the next versions of the AARC-BPA and it is a key topic in the Architecture work package of AARC2.

The **End Services Layer** contains the services users actually want to use. Access to these services is AAI-protected (possibly using different technologies, see Footnote 2). These services can range from simple web services, such as wikis or portals for accessing computing and storage resources, to non-web resources such as a login shell, an FTP transfer- or a workload management system. A notable change from AARC-BPA-2016 to AARC-BPA-2017 is the removal of the Token Translation Services from the End Services Layer. Credential translation or token translation can happen centrally and/or within a service, although the latter is outside of the scope of the AARC Blueprint Architecture.

AARC-BPA-2017 addresses most of the requirements that were still open in AARC-BPA-2017. Namely, in AARC-BPA-2017 we have provided:

- guidelines for expressing group membership and role information in a consistent manner across RIs/EIs;
- best practices for managing authorization, specifically targeting practices for community based authorization;
- guidelines for scalable attribute aggregation implementations;
- guidelines for the implementation of credential translation via token translation services;
- implementation scenarios and guidelines for credential delegation;
- implementation scenarios and guidelines for non-browser access;
- guidelines for implementing SAML authentication proxies for social media IdPs;
- use case scenarios for account linking and LoA elevation via step-up authentication.

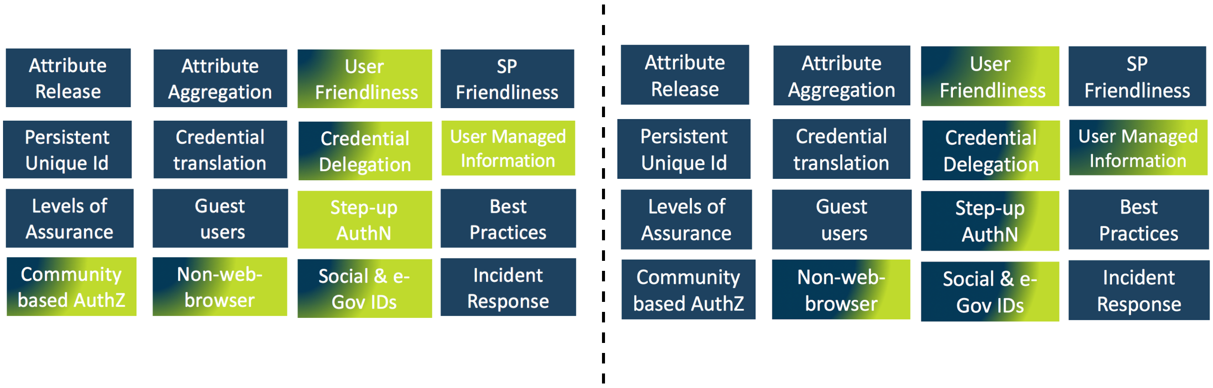

Figure 6: Progression of the requirements covered in AARC-BPA-2016 (left) to AARC-BPA-2017 (right).

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Blue are the requirements, which are already being addressed by the AARC. With blue-green colour, we have marked those requirements on which we have been partially addressed, but there is still work to be done. and with green colour, the requirements which will be addressed in the next iterations of the AARC Blueprint Architecture

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

# 3 Glossary

| | |
|---|---|
| **AA** | Attribute Authority |
| **AAI** | Authentication and Authorisation Infrastructure |
| **EGI** | European Grid Infrastructure |
| **FQDN** | Fully Qualified Domain Name |
| **IdP** | Identity Provider |
| **MACE** | Middleware Architecture Committee for Education |
| **NSS** | Namespace Specific String |
| **OIDCre** | OpenID Connect for Research and Education |
| **POSIX** | Portable Operating System Interface |
| **REFEDS** | Research and Education FEDerations group |
| **R&S** | Research and Scholarship |
| **SCIM** | System for Cross-domain Identity Management |
| **SP** | Service Provider |
| **URL** | Uniform Resource Locator |
| **URI** | Uniform Resource Identifier |
| **URN** | Uniform Resource Name |
| **VO** | Virtual Organization |
| **VOMS** | Virtual Organization Membership Service |
| **VOOT** | Virtual Organization Orthogonal Technology |

The research leading to these results has received funding from the European Community’s Horizon2020 Programme under Grant Agreement No. 653965 (AARC).