Template for West-Life Deliverables

Partner numbers

1	SCIENCE AND TECHNOLOGY FACILITIES COUNCIL	999980179	CO
2	STICHTING HET NEDERLANDS KANKER INSTITUUT-ANTONI VAN LEEUWENHOEK ZIEKENHUIS	999984738	BEN
3	EUROPEAN MOLECULAR BIOLOGY LABORATORY	999988230	BEN
4	Masarykova univerzita	999880657	BEN
5	AGENCIA ESTATAL CONSEJO SUPERIOR DE INVESTIGACIONES CIENTIFICAS	999991722	BEN
6	CONSORZIO INTERUNIVERSITARIO RISONANZE MAGNETICHE DI METALLO PROTEINE	999516810	BEN
7	INSTRUCT ACADEMIC SERVICES LIMITED	952324661	BEN
8	UNIVERSITEIT UTRECHT	999985805	BEN
9	LUNA SAS	942974540	BEN
10	ISTITUTO NAZIONALE DI FISICA NUCLEARE	999992789	BEN

Notes

Upload the deliverable report to

https://ec.europa.eu/research/participants/grants-app/reporting/DLV-675858

Check also for a related Milestone – if one is achieved, tick it and add a comment referring to the report.

Deliverable D4.2

Project Title:	World-wide E-infrastructure for structural biology	
Project Acronym:	West-Life	
Grant agreement no.:	**675858**	
Deliverable title:	Common security model design	
WP No.	4	
Lead Beneficiary:	Masaryk University	
WP Title	Operation and maintenance of the computing and data infrastructure	
Contractual delivery date:	Month 9	
Actual delivery date:	Month 9	
WP leader:	Ales Krenek	Partner: Masaryk University
Contributing partners:	Masaryk University, INFN	

Deliverable written by Daniel Kouřil, Marco Verlato, Chris Morris, Callum Smith, Pablo Conesa

1 Executive summary

Contemporary research infrastructures need to provide a sufficient level of security, which will protect users, service providers and operators of the infrastructures. In this document we focus on the area of authentication and authorization in West-Life since these functions are crucial for a secure design and also are exposed to end users and therefore significantly influence the users' perception of the system. The goal of the document is to overview current approaches and to provide a revision of the authentication and authorization mechanisms that will evolve during the West-Life project.

The design presented in the document harmonizes different approaches taken by existing components in order to achieve a single authentication and authorization infrastructure used thorough the project. The designed solution will enable easy integration of existing and new West-Life services and improve the usability of the systems for the end user.

The design utilizes mechanisms and principles that are commonly recognized in other research environments at the moment and provides a high level of interoperability. The solution will allow for integration with a unified authentication and authorization infrastructure if it emerges during the lifetime of West-Life.

2 Declaration of scope for security model

While security encompasses several areas, in this design document we focus mostly on functions that are most exposed to end users and hence they influence the overall perception of the environment by the community. Therefore, the main goal of the document is to review the existing authentication and authorization infrastructure (AAI) used in West-Life and provide its update based on current best practices and state of the art technology.

Other security functions are not covered by the document, however, they are often inherently provided as part of the environment. For instance, West-Life utilizes the security framework of the European Grid Infrastructure (EGI[1]) that defines security precautions for various layers of the infrastructures that are leveraged by West-Life.

We recognize that AAI is being widely discussed at the moment. Other infrastructures and projects try to revise their AAIs, pilot solutions that will verify new concepts or even move towards production utilization of newly established AAIs. We closely follow the development in projects and infrastructures leading AAI activities at the moment, like AARC[2], ELIXIR[3],

[1] http://www.egi.eu/
[2] https://aarc-project.eu/

Instruct[4], EGI and GEANT[5]. Several partners of West-Life are directly involved in the development performed in the scope of these projects and feed their experiences back to the project. We want to achieve a solution that is widely built on common principles so that it will enable easy cooperation with others. Anticipating further developments and unification of AAI solutions in other activities we want to arrive to a solution that could be easily adopted by or integrated with other AAIs.

The West-Life project continues to develop an infrastructure that is already established and used on regular basis for production work. Any significant revisions to the existing services must make sure they continue to work. Therefore, we plan to use components and technologies that are well tested and proved to work already.

2.1 Key areas of AAI

The domain of AAI is composed of three key areas that must be considered. We provide a brief introduction of the areas below.

2.1.1 Identity management

Thorough control over the life-cycle of a user in a virtual organization and/or infrastructure is crucial for secure operations of the infrastructure as well as the user's perception of it. An identity management system helps keep information about users consistent and current. A management of groups should provide a way to organize users and services into units with a dedicated purpose. The management should be flexible enough to facilitate users to organize their own groups and use them in the infrastructure.

2.1.2 Authentication

Authentication of users and servers establishes the digital identity of the entity. Depending on the actual profiles utilized by the authentication service in use, the authentication process can also mediate vetting of real identities of the user.

Traditionally, there are various aspects driven by the demands for usability, like support of Single Sign-On (SSO) or requirements to use multiple digital identifiers by a user. From the users' view, the most visible authentication service is the actual authentication method. There are various mechanisms that differ a lot and their selection can influence greatly the overall perception of the system by end users.

[3] https://www.elixir-europe.org/
[4] https://www.structuralbiology.eu/
[5] http://www.geant.org/

2.1.3 Authorization

Authorization decides about whether or not access to a service should be granted to a particular client. Authorization can be based on identifiers or additional attributes assigned to the client. Traditionally, attributes are maintained by a virtual organization or community where the user belongs to. Applications can also obtain attributes from other places, e.g. users' identity provider and another third-party attribute authorities.

3 AAI in current West-Life components

The West-Life environment encompasses various services and tools that differ in the way how AAI is supported. In this section we provide an overview of several selected tools to illustrate the current status.

3.1 AAI aspects in Cryo-EM applications

The Scipion Web Tools (SWT) currently do not implement any authentication mechanism at all. Access is completely open but privacy is guaranteed. However there are some security features in place and some others are planned to be implemented.

Users can create a project and upload their own data. The way SWT protects a project and its data from being accessed by a different user is implemented by associating a unique URL with the project, which is constructed using a cryptographic hash that is hard to reverse. Anyone with that URL will be able to access the project with exactly the same privileges. There is no search interface to find a project, so no user can discover someone else's projects. The main drawback of this approach is that users must bookmark or somehow keep that URL somewhere in order to access the project later. Only server folders within the project can be explored, not allowing the users to explore folders containing data belonging to other users.

Another relevant aspect is how to restrict the use of the resources to avoid their abuse. Currently, all the processes that can be executed on the server have predefined, read only, values for parameters defining the number of threads and MPI specifications, not allowing any user to change them. These processes are run on a local cluster managed by the Slurm batch system.

There are a number of benefits expected for the SWT from a West-Life wide AAI:

- Identification: with an authentication mechanism in place we will free users from bookmarking their projects and present a list of their own projects once they are logged in the portal. We will also benefit from knowing their identities, which is not available in the current implementation. We could also communicate with them via email (if provided) to notify about status of the processes.

- Access level: with user accounts in place we could implement different access levels to projects, allowing external users to access projects in read-only mode, or with certain editing privileges defined by the owner, allowing a certain degree of collaboration.

- Resources allocation: policies based on user identity or group membership can be adopted to balance and optimise the use of computational resources.

- Distributed data access: interoperable AAI mechanisms allow data owners to grant access to user data stored in external infrastructures (e.g. EUDAT[6]).

3.2 AAI aspects in WeNMR infrastructure

The WeNMR infrastructure [8] provides a platform integrating and streamlining the computational approaches necessary for NMR and SAXS data analysis and structural modelling. The platform consists of a VRC portal[7] from which a number of application portals, developed and hosted at different NMR or SAXS laboratories, can be accessed through a Single Sign On (SSO) mechanism based on the SAML protocol (via a SimpleSAMLphp implementation).

Access rights to the application portals were in fact guaranteed by requiring:

- a registration to WeNMR VRC portal, that enables SSO and accounting facility for all the associated WeNMR application portals. Authentication takes place via username/password or via the SURFconext identity federation (for Dutch users only).

- a X509 personal certificate issued by a IGTF accredited CA loaded into the users' browser

- the personal X509 Certificate Distinguished Name (DN) had to be registered with the enmr.eu VO

The X509 certificate requirements are only needed to getting access to CPU-intensive services offered by WeNMR such as the CS-Rosetta and AMPS-NMR portals.

[6] https://www.eudat.eu/
[7] http://www.wenmr.eu/

Behind the application portals a business logic software component performs the needed operations to allow job submission over the EGI HTC platform. In fact, the large number of jobs to be submitted and monitored by most of the application portals strongly required automation. For what concerns AAI, the current version of the HADDOCK, CS-Rosetta, AMPS-NMR and UNIO portals job submission is implemented making use of X509 robot certificate issued by the NIKHEF and INFN certification authorities, according to the EGI VO Portal Policy document [2] published by the Security Policy Group of EGI. These portals (which are currently linked also from the West-Life/Services page[8]) are qualified as a "parametric portal" according to the policy document. Robot certificate's names are registered with the enmr.eu virtual organizations, under a VO group that identifies the application. The VOMS servers supporting the enmr.eu VO are hosted at INFN in a high availability mode (voms2.cnaf.infn.it is the master server, voms-02.pd.infn.it is the backup server). VOMS-enabled proxy certifications with group information are periodically created starting from the robot certificates and used to submit grid jobs on behalf of the user.

A custom user tracking system was implemented by portal developers in order to keep track of applications usage by each user. To maintain privacy, usual grid accounting tools will see only the aggregate application usage, based on the VO group information as shown in [10], while the portal administrator will be always able to see the usage records of each user.

WeNMR project developed a SSOXS (Single Sign On for eXternal Services) module [9] that extends the Drupal content management system to serve as an authentication and central user management entity and accounting portal. Features of the module include:

- Administrators can register services for which to enable SSO and users can register for services on their user account page
- SSO credentials are stored in a secure external database accessible to the service
- Powerful accounting functionality to track service usage and performance

The SSOXS module is available from the Drupal module repository[9].

3.3 AAI aspects in CCP4 infrastructure

CCP4 is a suite of software applications used to solve and validate structures by X-ray crystallography. The applications are available through an online service[10]. Users register to

[8] https://portal.west-life.eu/services/
[9] https://www.drupal.org/sandbox/mvdijk/2123977

it via username/password. The possibility of using ORCID to link different identities together is under investigation.

A second service, CCP4-DaAS, is using the STFC SAML based IdP service for authentication and authorisation for the initial implementation. This is to facilitate users accessing data collected at Diamond Light Source and residing in ICAT, and submitting jobs to the SCARF cluster. The possibility to use other credentials is under investigation, and may be adopted once the prototyping phase is completed. The STFC IdP has been proved to interoperate with Umbrella.

CCP4 software is also available for download, and is free for academic use provided that the user accepts the license. As a proof of concept, the STFC West-Life team has built a VM image containing the CCP4 suite, and is discussing with CCP4 the licensing arrangements.

3.4 AAI aspects of ARIA (Instruct) and integration to EUDAT

ARIA is a collection of services produced by Instruct that offers tools to biological and medical sciences (BMS) facilities, infrastructures and user communities as a SaaS model. Behind these tools is an AAI designed to support the needs of the BMS user community and has a series of identity management tools integrated to the core services to allow for cross-service interoperability. Key to this is a central dynamic groups system which can be manipulated to assign group ownership and rights that are administered centrally within ARIA and distributed to varying services that utilise the ARIA IdP. The level of release is controlled by the Shibboleth IdP v3 software which also provides the SAML authentication workflows and transactions for the service. This allows the identity service to have the intricate control of user rights management and group management whilst being able to automate as much of the management of these rights through ARIA. Additionally, high-level controls allow for manual intervention over any aspect of the rights management.

Multiple identity providers are connected to the ARIA AAI to allow maximum user choice as to where user credentials are stored and provided from. Internally software for account credential linking has been developed to allow multiple identity sources to be resolved against a single internal ARIA account similar to that of WeNMR SSOXS.

Through engagement in West-Life cross-AAI links have been enabled between ARIA and EUDAT to allow existing users of ARIA AAI to auto-provision accounts within EUDAT. All prerequisite registration data is provided through the SAML attribute channel directly from

[10] http://www.ccp4.ac.uk/ccp4online

ARIA to EUDAT. Users experience an entirely transparent and fast transition from their existing community to the services and tools offered through EUDAT.

4 Proposed AAI in West-Life

Having taken into account current state of several West-Life tools articulated in the previous section and trying to achieve a harmonized and interoperable AAI solution we have reviewed the mechanisms that are used nowadays and designed a new generation of AAI for West-Life, which is presented in this section.

The revised architecture follows recent developments in other similar infrastructures and is based heavily on outcomes of the AARC project that strives to harmonize AAI developments in the research and educational area. The architecture also responds to the requirements that we reported in a submission to the AARC user survey document [1].

We based our approach on the current state of the AARC Blueprint Architecture [4], which provides a reference architecture to help design proper AAI solutions. The design is interoperable and based on current approaches, which ensures compatibility and extensibility with further development.

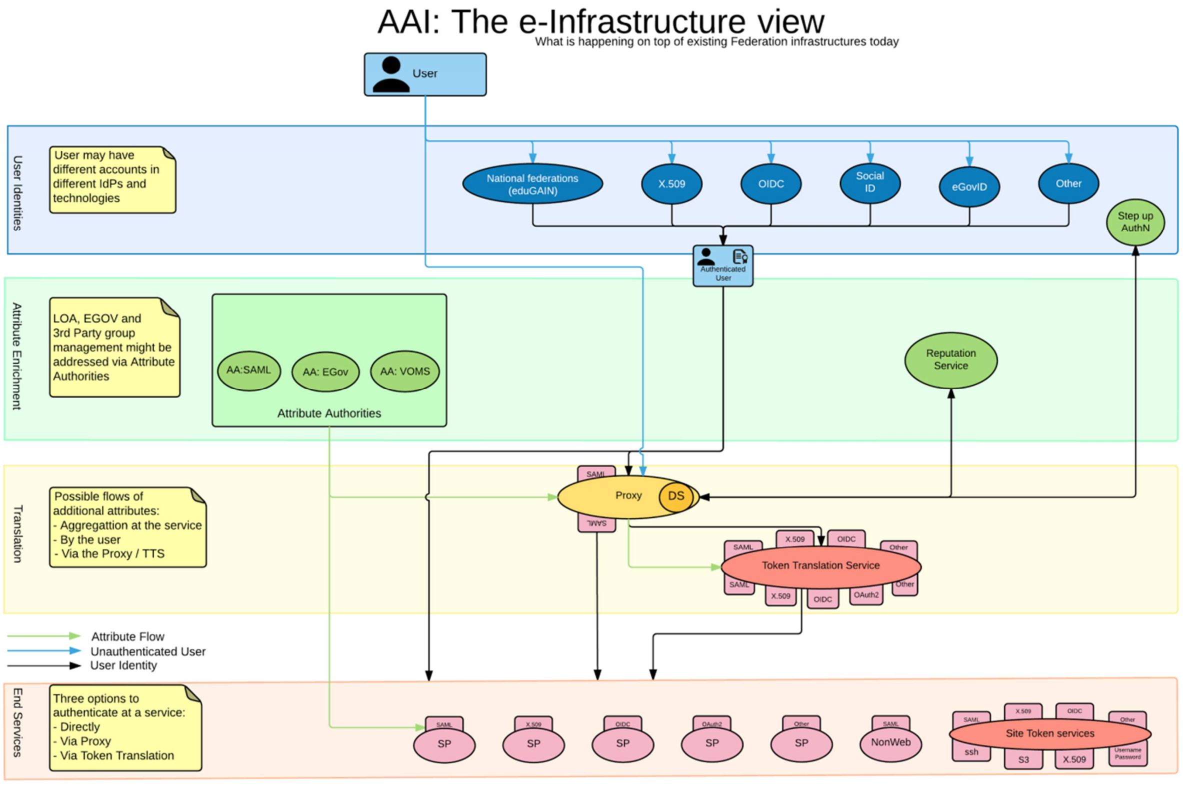

Figure 1: AAI reference architecture schema by AARC

The generic AARC architecture is depicted in Figure 1, and is composed of four layers that are described below.

The **User Identities Layer** contains services for identification and authentication of users. There are multiple providers of identities even for a single user, which may use different authentication technologies.

The **Attribute Enrichment Layer** contains services related to managing and providing information about users that are shared as additional attributes. Typically, they provide additional information about the user. While certain user-related information can be provided by the authentication provider, it is not always sufficient and extended attributes are maintained elsewhere, such as memberships in virtual organizations, etc. The information is used for authorization decisions taken by end services.

The **Translation Layer** addresses the requirement for supporting multiple authentication technologies. Services on the layer translate credentials from one format to another to provide a bridge between different mechanisms. Other services transform trust relationship linked to authentication decisions made by user's identity providers.

The **End Services Layer** contains the services users want to use and that control access to their resources based on information provided from the other layers.

From the architecture point of view, there is a significant difference among the layers given by the way how they function. The first three layers are typically deployed and operated in a controlled manner, usually by institutions or large user communities. The number of their instances is limited and they are carefully hardened to avoid compromise. End services, on the other hand, are operated in a large number of instances and are not controlled globally by the infrastructure provider.

The AARC reference architecture outlines fundamental domains that should be taken into account when an AAI is considered. In the rest of this section we provide a more detailed discussion of the layers with respect to the West-Life environment. In addition to the reference architecture, we consider challenges listed in the AARC Blueprint document as well as our experiences gained during establishments of AAIs in other life science activities (namely ELIXIR and Instruct) or infrastructures (EGI, GEANT). The design takes into account existing services and is provided in a way that enables their smooth transition to new AAI functions.

4.1 Identity management

In addition to the layers described by the AARC model, we explicitly consider management of users' identities and necessary functions for it. The goal of an identity management (IdM) system is to provide processes for handling users and their identities in the whole system that used by them. With respect to the AARC reference model, the IdM system provides the basis that delivers data to services at various levels.

A central identity management (IdM) system for the whole West-Life environment will be maintained. IdM must be provided as a single logical service that is operated on the infrastructure level. The West-Life IdM will integrate several types of functions:

- A full control of users' life-cycle, i.e. registering of new users, renewal their registrations, suspension of user accounts and ceasing their registrations (on expiration or removal).
- Group management allowing for assignment of maintained identities (users) to groups. Users should be allowed to maintain their own groups and control its membership.
- Account linking to connect multiple identities to a single user records so that the user is always recognized equally, regardless the authentication method or credentials used.
- (De-)provisioning of user records that makes it possible to propagate information about users to end services when needed (e.g. to populate user accounts).
- Management of various levels of assurance for user records and attributes maintained.
- Support of the concept of a virtual organization that manages all users in a single body.

IdM components can be provided as separate services integrated together or standalone software providing all the functions. The user, however, must have a unified interface to work with the system.

We have to make sure the IdM system handles easily the existing bases of users, namely those from WeNMR and Instruct.

4.2 Authentication and user identities

Authentication is a basic function using which services establish the identification of their clients and vice versa. For the environment of West-Life we use the concept of an identity provider (IdP) that provides the authentication service and asserts the digital identity of users that belongs to it. Besides authentication decision, IdP also provides an identifier that is connected to the authenticated user. At the moment several technologies exist to implement the IdP (see e.g. [7]). For the domain of web services, protocols based on SAML have become common, especially in R&D area. IdPs provided by large social networks are usually based on the OpenId Connect protocol, which provides a similar user perception of the authentication process. Similar technologies are expected to be used to provide government eID or eIDAS identifiers. Systems that evolved from Grid technologies utilize X.509 certificates that are issued by identity providers (certificate authorities) according to rules by the Interoperable Global Trust Federation (IGTF).

We plan to enable support for external identity providers, such as Google, LinkedIn, and ORCID. We will also support identity providers that are utilized heavily by the current user base, namely to support existing WeNMR and Instruct identities.

Identity providers provide a different level of identity vetting of the user. There are different levels of assurance recognized and required, mostly based on requirements of end services. It is important to realize that authentication process verifies the claimed identifier but does not necessarily provide vetting of identities. For instance, IdP provided by large social networks like Google and Facebook provide quite strong level of credential verification (including e.g. two-factor protocols) but does not guarantee anything about the real identity of the user, which renders the identities untraceable to real-world legal identities. Since some services will require traceability of users, additional information must be introduced to the West-Life IdM, which will help provide additional information about the user.

In order to obtain higher levels of verification, we will enable employing of government controlled identification systems once they become usable (driven by the EU eIDAS regulation and piloted in STORK activities)

The technologies differ in various aspects and there is no a single authentication technology that would cover all requirements and provide access to all services in contemporary systems. The West-Life AAI has therefore to be able to interoperate with different authentication technologies and also allow users to use multiple identifiers (each linked to a different credential name as determined by the technology and the profile of the user's IdP).

West-Life will be an identity consumer and reuse existing identities; we do not intend to be an identity provider. We will require users to have an account at one or more external trusted Identity Provider (e.g. Instruct) so that authentication can be delegated to them. The potential end users are structural biologists in Europe. Most are already recognized by IdPs of Instruct, WeNMR, eduGAIN, and ORCID.

In order to get access to the infrastructure, users have to be provided with a user-friendly mechanism that will follow the Single Sign-In principle. The mechanism must protect the credentials sufficiently.

To facilitate complex workflow where several peers are involved in the processing chain, it is often necessary to support delegation of credentials so that services in the chain could work on behalf of the user that initiated the operation. Current West-Life service use dedicated service account to support such workflows (e.g. web portals submitting grid jobs using robot certificates) and therefore the need for credential delegation is not urgent. We will examine the actual needs in the new types of service and possibly evaluate existing approaches to support credential delegation.

4.3 Attribute management

Attributes provide additional information about the user, which can be used for authorization decisions, personalization of service interfaces, etc. Attributes can be managed at various places. Identity providers typically maintain attributes that are available to the users' home organization, like the name, email address. There are, however attributes that need to maintain elsewhere, like information provided by the virtual organization or community of the user. In order to maintain these attributes it is necessary to establish a service (often called an Attribute authority) that will maintain them and make them available to relying parties. A typical example is VOMS that releases information about group membership and VO-related information in the context of the WeNMR VO in EGI.

There is no limit on the number of attribute authorities and the West-Life AAI must be able to cope with attributes issued by multiple authorities for a single user and transport them to end services in an aggregated form. West-Life should be also able to maintain a set of attributes that is under control of the community, which will make it possible to maintain e.g. persistent unique identifiers assigned to every single user.

Attributes provides a powerful mechanism to express descriptions related to a user and can bear arbitrary pieces of information. However, similarly to identity vetting or authentication, the level of assurance must be checked. We will follow the guidelines set by AARC and will monitor the development in the area to integrate further recommendations.

4.4 Token translations

As mentioned before, West-Life users utilize services that require authentication using X.509 certificates, such as job submission to the Grid or management of virtual machines in the EGI Federated Cloud. X.509 credentials are therefore needed to access these services but they also were recognized as posing a significant barrier for usability. In order to facilitate access to X.509-based services without forcing user to maintain certificates, several approaches appeared.

A solution that is often in use at the moment employs a dedicated service that interacts with the Grid (or other X.509-based services) on behalf of their user. The service uses a *Robot certificate* to authenticate itself so that they act as the client for other services. Users authenticate by other means to the service, either by a federated mechanism or using a local credential managed by the service. This approach is used by web portals in the environment of West-Life.

There are also other ways how easier access to of X.509 credentials can be facilitated for end users, including X.509 proxy repositories and on-line certification authorities. For the needs of the West-Life environment it is required to examine a technology that not only makes credentials retrieval easy but also hides the management of the credentials so that end-users are not forced to secure the private keys themselves.

4.5 End services

End services provide access to resources for the user to perform their work. They are responsible for controlling access to the resources for which they need enough information from other layers of the AARC architecture. They rely on identities provided by the authentication services and possibly on additional attributes that are assigned to the user accessing the service. Different services may have different requirements on the range of the attribute and their levels of assurance. Some services accept lower levels while other services are more critical and require the highest levels of confidence. It is always responsibility of the service provider to specify requested levels. The West-Life AAI will provide sufficient methods to deliver the information needed.

Access control rules in West-Life will be based on group and/or roles of users and we will try to avoid listing particular user identifiers in ACLs. The extent to which this type of access control can be specified depends on particular applications. Application running in “containers” can leverage from access control that is provided by the container (like the Apache web server). If access control decision should be performed on the level of end service, it should be able to process information about user’s membership in groups or roles.

Given the variety of services, it is necessary to ensure sufficient flexibility so that service can be accessed using the security mechanisms they natively support. West-Life services are typically provided as web portals where authentication is well understood and several proven mechanisms supported. However, we expect also other types of services to emerge that will not support the standard access with a web browser, for which the AAI will have to function as well.

West-Life utilizes the concept of service accounts (authenticated using e.g. robot certificates) that are used for web portals. The concept will be used and possibly extended for other types of end services. For instance, WP6 is expected to link cloud services identities including Dropbox, Google, Amazon into one West-Life account for virtual folders.

5 Implementation of revised AAI in West-Life

The previous section summarized functions of the AAI layers that the West-Life AAI should deliver. We foresee that the new AAI functions will be introduced gradually and in several iterations so that their impact on infrastructure is properly assessed. This section describes the first set of services to be introduced and evaluated.

The suggested AAI components and their integration with the infrastructure stem from experiences with pilot AAIs that were verified in the context of activities performed in the environments of ELIXIR [5] and Instruct. This experience should ensure interoperability with other solutions and also the possibility to hand over the operations of verified solutions where possible.

An active development in the area is on-going, both in communities and infrastructures, our solution should be able to examine possibilities that will emerge and possibly integrate them. Figure 2 shows a scheme of an updated AAI for West-Life. Individual services of the schema are further elaborated in the rest of the section.

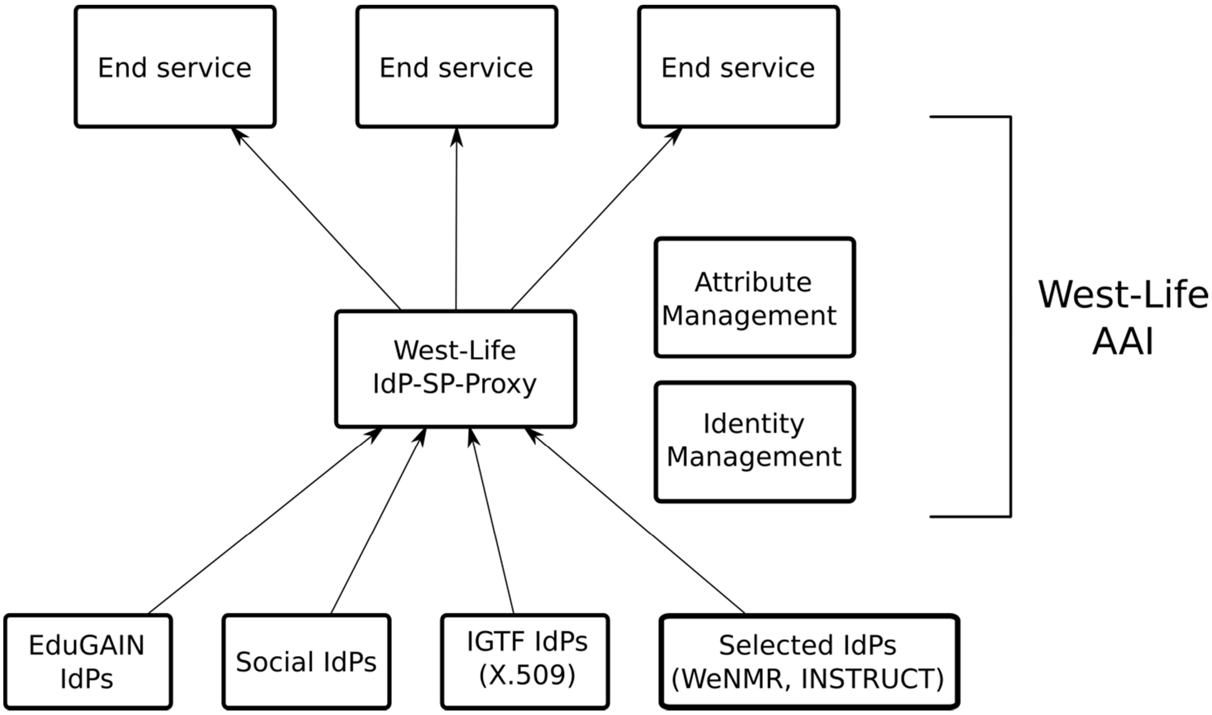

Figure 2: Revised West-Life AAI

5.1 Identity management

Current implementation is based on the WeNMR SSO module (see Sec. 3.2). We plan to exchange the module with a solution that is tested in other similar environments. There are several tools available, as overviewed e.g. by AARC. We will start with evaluation of Perun[11] [6] that is used for similar purposes in ELIXIR and EGI.

User’s lifecycle will remain unchanged, including users’ registration. Integration of IdM system with other services will provide a single interface to the users and they will not be required to directly register with other services that play a role in the management of identities (e.g. VOMS). No additional credentials will be required for the registration, which must be easy to perform for user assessed by recognized IdPs.

For the level of assurance of identity vetting in the initial phase we will continue to rely on X.509-base trust model maintained by the IGTF and will allow users to “elevate” their LoA by presenting their X.509 credentials from a CA accredited by a proper IGTF profile.

[11] https://perun.cesnet.cz/

5.2 Authentication services

As presented in the previous sections, the West-Life AAI will be open to various different IdP providing users authentication. In order to ease the introduction of IdPs and decrease the administrative overheads, we will employ the IdP-SP-Proxy concept.

An IdP-SP-Proxy refers to a common pattern used in contemporary AAIs, which provides a bridge between internal structures of the infrastructures and external services. An IdP-SP-Proxy can be deemed as a translation service that transforms not only authentication format but also helps address several crucial concerns and challenges with AAI.

Several possibilities exist that provide the functionality needed. We will start with suggestions articulated by the AARC technology overview [7] (for instance with OpenConext by SURFnet).

We will continue to support service identities for web portal together with robot X.509 certificates. Developments of solutions enabling translation of users' X.509 credential (like e.g. CI-Logon pilots in EGI and ELIXIR) will be followed and their utility evaluated once more experiences with their piloting is gained.

5.3 Attribute management

A basic set of attributes will stay maintained by the West-Life services. The VOMS service necessary for EGI service will be maintained as a backend service so that users will not be required to directly interact with it.

With different IdPs it is necessary to address different identifiers of a single user. The West-Life IdP-SP-Proxy will be able to link multiple identifiers to a single user identity so that the end service will have a consistent naming of the user. It will also be possible for the proxy to assign a persistent identifier to the users that will be solely maintained by West-Life independently on external IdPs.

5.4 Integrating AAI with end services

Providers of end services have to specify requirements on authentication and access control. AAI will provide sufficient information about the user so that the service could decide if access is granted or not for a given user. Information will be passed as attributes.

In terms of protocols, SAML will be used to interface existing IdPs. For new services we plan to support OIDC for authentication and attribute exchange however SAML will still be supported as a mechanism for services that exist already or that are can support SAML easily.

5.4.1 Web front-ends

Most existing services are provided as web-based portals where authentication will be based on a web SSO system. While SAML-based solutions are very common at the moment, we will also add support for OIDC, which provides the same level of user experience but is much easier to integrate with the service. The West-Life-IdP-SP proxy will be used as the authentication provider from the point of the view of the services. Authorization will be based on attributes released and/or passed by the West-Life-IdP-SP proxy as well.

5.4.2 Cloud management

West-Life plans on utilization of IaaS cloud facilities for their users. Since cloud resources will be provided by a third-party provider, the needs for AAI will be dictated externally. For the EGI Federated Cloud X.509 proxy certificates are required for the user to maintain their virtual machines; we plan to utilize a similar approach to handle computing jobs and employ robot certificates for service identities. The VOMS attribute needed to pass information about the VO in EGI will be populated by the established VOMS server. We plan to adapt the existing VO registration process so that the VOMS server is populated directly from the West-Life IdM and users do not need to explicitly register with the VOMS server.

5.4.3 Accessing data storages

For management of large data volumes we plan to leverage as much as possible solutions based on EUDAT, like the B2SHARE service. AAI services would either utilize service identities or recent development that achieved integration of EUDAT and ARIA (c.f. 3.4).

For direct POSIX-like access to file systems in EUDAT, B2DROP makes it possible to mount a remote file system using the WebDAV protocol. Support for AAI in B2DROP is undergoing revisions by EUDAT and we will strive to enable the West-Life AAI once more details are available.

Isolated virtual clusters that will need to share data among the nodes could use any distributed file system that is supported by the application with authentication being provided by ad-hoc services (like internal Kerberos deployment).

5.4.4 Virtual folders

In order to provide an aggregated view of all scattered data of a user the West-Life will provide a *virtual folder* that can be mounted and unifies access to files stored on different places. The access to majority of services will be provided via the HTTP protocol and we will need to closely examine the requirements of individual providers. The current development seems to use Apache web server, with reverse proxies to other application services based on context.

5.5 Towards an interoperable and sustainable AAI

The West-Life AAI services will be a core component of the whole infrastructure and therefore it is crucial to design and establish it in a way that will enable continuity after the project is finished.

The West-Life users originate from the area of biological and medical sciences (BMS) in Europe that use also services provided also by other projects and initiatives. It is important that the users have smooth access to the service and can have a unified access to all the services available in the domain. The West-Life AAI should therefore be interoperable to cover the needs.

The need for an interoperable and sustainable AAI is recognized by other projects and activities. West-Life partners have already enabled discussions between Instruct and EUDAT to this end. There are also discussions on-going as part of the CORBEL project[12], which provides a forum for AAI collaboration in research infrastructures in biological and medical sciences. Sustainability of AAI solutions is also being discussed in the context of the AARC project.

West-Life partners are actively involved in several these activities. We leverage direct involvement of our partners in ELIXIR that piloted a reference AAI model that is being verified at the moment and is moving to production in starting from September 2016. The ELIXIR AAI could also serve as a prototype for future BMS-wide AAI. ELIXIR produced a strategy for the adoption and use of an AAI [3], which outlines basic principles for a sustainable and user-friendly AAI, enabling a close cooperation within the BMS domain. The strategy envisions collaborations with e-Infrastructures (e.g. GEANT and EGI) that will provide AAI services as part of the portfolio they offer to their user communities. Additionally through the newly funded project AARC2 strategies looking to explore cross-AAI connections are being explored by Instruct, EUDAT and other AAIs. West-Life will closely

[12] http://www.corbel-project.eu/

follow the development to be able to adapt and integrate results to achieve a long-term sustainable solution. The West-Life AAI was designed in a way that should ensure interoperability with AAI solutions that are expected to emerge. During the project we will evaluate possible approaches and prepare the set of services that can be offloaded to an external AAI provider (like an e-Infrastructure) once it is established for the BMS community.

References cited

[1] C. Kanellopoulos, N. Liampotis, N. van Dijk, P. Solagna (editors). Analysis of user community and service provider requirements. *AARC Deliverable DJRA1.1*. 2015. Available from https://aarc-project.eu/wp-content/uploads/2015/10/AARC-DJRA1.1.pdf.

[2] EGI Security Policy Group. VO Portal Policy. 2010. Available from https://documents.egi.eu/document/80.

[3] ELIXIR project. Strategy for the adoption and use of an AAI within ELIXIR. 2016. Available from https://docs.google.com/document/d/1cJ3mR8lqfZKRMvSFalSmPbqd1OPU-L6YcUFIRnh1rhQ

[4] M. Hardt et al. First draft of the Blueprint Architecture. *AARC Milestone MJRA1.4* (work in progress). 2016. Available from https://aarc-project.eu/aarc-draft-blueprint-architecture-available-for-comments/.

[5] M. Linden (editor). ELIXIR AAI – Requirements and Design. 2015. Available from https://docs.google.com/document/d/1CMY1np3GyvPD8LcKvXljXcRO04V2zu3n_Jcg19jgNOw/edit.

[6] M. Procházka, S. Licehammer, L. Matyska. Perun—Modern approach for user and service management. In *IST-Africa Conference Proceedings*, IEEE. 2014.

[7] P. Solagna et al. Existing AAI and available technologies for federated access. *AARC Milestone MJRA1.1*. 2016. Available from https://aarc-project.eu/wp-content/uploads/2016/01/MJRA1.1-Existing-AAI-and-available-technologies.pdf.

[8] T. Wassenaar et al. WeNMR: Structural Biology on the Grid. *Journal of Grid Computing*. December 2012, Volume 10, Issue 4, pp 743-767.

[9] WeNMR releases its Drupal module for Single Sign On for eXternal Services (SSOXS). *WeNMR News*. 2013. Available from https://www.wenmr.eu/wenmr/wenmr-sso-module.

[10] West-Life project. Inventory of available resources and testbed setup*. West-Life Milestone M11*. 2016. Available from http://internal-wiki.west-life.eu/w/index.php?title=M11.