Date: 22 July 2014
Author: Licia Florio <florio@terena.org>

Report from the AAI Workshop and consultation meeting

Introduction

On April 2nd 2014, the EC hosted an AAI Workshop, as a follow-up of the AAA Study on Supporting Scientific Data and as a stakeholders consultation for the Work Programme 2016-2017 of the Horizon2020. This document reports on the outcome of the workshop.

The workshop programme was coordinated by TERENA and sponsored by the GÉANT project. The programme and list of attendees are attached to this report as well as online at: http://www.terena.org/activities/AAI-Workshop/agenda.html

Table of Contents

1. Executive Summary and Recommendations for WP16-17 2
1. Recommendations for WP16-17 3
2. Workshop Report 4
1. Policy Setting 4
2. Developments in the AAI area 4
3. AAI Community Requirements 5
4. Summary of the Break out Sessions 7
Annex I – e-Infrastructures and the AAIs they offer 9
Annex II – Workshop Agenda 11
Annex III – List of Attendees 12

1. Executive Summary and Recommendations for WP16-17

This document reports on the discussion and presentations held during AAI workshop in April 2014. The event besides offering a meeting opportunity, provided also a consultation mechanism to inform the EC on specific actions for the Work Programme 2016-2017.

The workshop explored a comprehensive set of initiatives from different angles: identity providers and federations, content/community providers and e-infrastructure providers. This was combined with real life demonstration of (inter)federation framework (OpenConext and eduGAIN see more in chapter 2.2).

It was widely recognised in the workshop that identity plays a very important role in the way resources are being accessed. Identity should in fact be regarded as a core element of e-infrastructures of equal importance as network connectivity. This extends beyond access to services to a more holistic approach to managing identities, including the use of personal identifiers. Inter-operability of different user credentials is a high priority for all e-Infrastructures for establishing an eco-system where users, identity and attribute providers and resources providers smoothly and securely meet.

Work to enhance AAI and inter-federation frameworks as well as to address harmonisation of e-Infrastructures is already covered by a number of initiatives notably within the GÉANT community. Further activities are planned within the context of the Work Programme 2014-2105 (i.e. for the GN4 Framework Partnership and in the response to the E-INFRA-7 call).

Thanks to a number of parallel working sessions, the workshop helped establish a list of requirements and needs for those future developments:

- Need for training for institutions on the main principles of the EC data protection and the GÉANT Data Protection Code of Conduct (used by eduGAIN);
- Need to ramp up skills and expertise in institutions on AAI technologies;
- Need to position clearly the role of various e-Infrastructures in the end-to-end delivery of federated identity services and their added-value;
- Continue the work on group management/attribute providers, by integrating existing technical solutions in the various e-Infrastructures;
- The inter-operability between the R&E sectors with e-Government is in a very early stage. Support should be given to harmonise e-ID deployments in the various member states;
- Identify requirements and procedures needed by service providers to define a scalable and inclusive LoA framework that can be supported by institutions;
- Strengthen the work to offer ready-to-use solutions (i.e. Moonshot, OpenId Connect, OAuth2) for classic non-Web SSO as well as to support growing mobility access;
- Support the deployment of federations, by providing tools to create IdPs in institutions with limited manpower or know-how, and by supporting guest users;
- Define a common e-Infrastructure security and policy framework;
- Ensure smooth interoperability between AAI and Digital Identifiers;
- Define and operate a testing environment as well as associated tools for ensuring robustness and scalability of AAI protocols and policies.

1. Recommendations for WP16-17

During the workshop a number of topics were identified that have a long term perspective. The following **specific actions for WP16-17** are proposed mostly to follow up on WP2014-15 work but also looking at new areas:

- Ensure that an eco-system of e-infrastructure operators fully address important European research developments such as: Elixir, Copernicus or Human Brain Project where AAI plays a significant role;
- Increasing globalisation of education requires higher harmonisation of attributes to express roles, education levels and achievements;
- Stimulate open and creative usage/evolution of AAI in various form: open call, hackathon, competition, market place etc.;
- Stimulate at European level, usage of existing AAIs by making content accessed via these AAI;
- Operate the enabling technologies (notably accounting) for supporting various business models such as pre-paid (licence credit, one-off…) and post-paid (per usage, per licence);
- Create an acceptable LoA framework that can be supported by e-Infrastructure services and institutions;
- Continue the test and deployment of inter-operability between R&E and e-government communities and ensure convergence of national AAI roadmaps;
- Continue to engage in the standardisation effort;
- Strengthen the testing environment to facilitate the introduction of new features in the existing e-Infrastructures;
- Pursue skills ramp up and community management in a form of trainings, career path, and exchange programmes.

2. Workshop Report

1. Policy Setting

Carlos Morais Pires (EC) and Jean-Luc Dorel (EC) setup the policy context by recalling the main output of the AAI Study, the currently open call topic for AAI (EINFRA-7) and the review of the progress made over the last year with regards to the AAI study main recommendations and set the expectation for the workshop:

- Facilitate communication among different groups (IdP, SP, IdP+SP, users, federations, interfederation etc.);
- Coordination among different efforts
- Stakeholders consultation for WP16-17

Reviewing the progress, the following was highlighted:

- On the technical side, work is still on-going to address Single Sign On for non-web applications as well on the authorisation side. Both these areas are rather complex and even if progresses have been made, there is still need for further work.
- On the policy side, it was noted that there is a variety of initiatives where work on this area is carried out, such as the GÉANT project, EGI, IGTF, REFEDS, e-IRG and so on. Whilst the diversity of efforts is good and there is no intention to create a monolithic block, it is important to understand roles, interfaces between the players and to ensure collaboration and coordination among them. The EINFRA-7 call in the Horizon 2020 offers an opportunity to reinforce existing collaboration between e-Infrastructures, REFEDS, FIM4R and user-communities.

2. Developments in the AAI area

The latest developments in the AAI area were presented by the representatives of the different e-Infrastructures, namely GÉANT (eduGAIN), EGI, EUDAT and STORK2. Each of these AAIs was developed to address specific use-cases. New use-cases keep emerging and the user-base keeps growing, making each AAI a dynamic environment that evolves to offer new features. It is therefore important to continue the development of these AAIs further and to place it in the existing e-Infrastructures that are as close as possible to the communities from which the use-case originates.

The table below highlights the main features of the presented AAI, which communities they serve, the underlying technology and whether the infrastructure is production.

Name and Type of AAI	Community	Features and Status
eduGAIN (SAML2)	NRENs/campuses	Federated access for Web applications Production
EGI (x.509 certificates mostly)	eScience	Access to Grid resources (incluiding support for non-Web applications). Production
EUDAT (a mix of technologies)	Data Centres	Access to data management/storage resources (web and non-web) Under development
STORK2 (PKI encoded in eIDs and STORK SAML profile)	Government and citizens	Access to public services for business Pilot

Table 1: AAIs main features

With more maturity reached by various AAIs, it is crucial to look at horizontal actions to ensure the development of interfaces among them and that harmonisation and interoperability are achieved. The EINFRA-7 call offers a good example of what can be done; this type of work should be continued in the WP16-17.

Annex I reports more in-depth on the presentations on the various AAIs held during the workshop.

3. AAI Community Requirements

A number of different user communities attended the workshops and reported on their requirements. The table below shows the community requirements and how they are met at the moment.
As the table shows, a number of requirements are common to several communities: guest IdPs, attribute release on the IdP side, level of assurance, privacy policy and SSO for non-web applications.

Community	Requirements	Requirements Met	Requirements not Met
DARIAH (Digital Research Infrastructure for the Arts and Humanities)	- easy to use - ideally allow researchers to use the same credential in any 'academic' context. - SSO to all DARIA resources, tools, and services use same credential in any academic context - authorisation granularity	- guest IdPs and attribute providers operated by DARIAH - integration into eduGAIN works via DFN-AAI	- eduGAIN has too little outreach, not every institution signs federation contracts - not every IdP releases necessary attributes; - technologies for non-web-based access only "almost there”
IGTF (International Grid Trust Federation)	- different level of Assurance - authorisation controlled by the community - non-Web SSO - attributes release - federations (and IdPs) to work on a collaborative security and policy framework - support for citisen researcher	- most of technologies are there, but effort is needed to stich them together - Code of Conduct is a big step	- LoA not supported properly - SSO for non-Web not really available - IdPs still not releasing attributes
CLARIN (Common Language Resources and Technology Infrastructure) and DASISH	- discovery - guest IdPs - SAML/OAuth2 bridge for trust delegation	By providing: - a federations of service providers - guest IdPs - International discovery service	- Broad coverage of federations - attributes not being released predictably
ELIXIR (European Life Science Infrastructure for Biological Information)	- Level of Assurance - uniform service (i.e. opt-in, attributes, etc.;) - good coverage of R&E federations - SSO also to cloud resources	- eduGAIN is a good start	

CERN	- well-defined framework to ensure sufficient trust and security among the different IdPs; - attribute definition and release; - authorisation handled by the community; - Level of assurance Web and non-web SSO	- key technical building block identified - key forums and groups identified	- security policy framework not clearly defined - operational security - well defined LoA - ECP profile
ESA	- enable ESA SSO to join existing federations; - establish a space identity federation framework to simply the cooperation among space partners	- specifications for some of the building blocks already available	- not all ESA partners are ready to move to federated access

Table 2: Community requirements

4. Summary of the Break out Sessions

A break out session followed to discuss four different topics:

- Attribute Release:
- Level of Assurance
- Collaboration
- SSO for non-web

A short summary is reported below.

The **attribute release** group discussed about the issues that revolve about releasing attributes as well as the quality of the attributes. eduGAIN is often blamed for not working properly, but in reality IdPs of the federations participating in eduGAIN do not share attributes in a consistent way. This problems becomes more acute when IdPs are asked to share users' personal information to a service the is not operated in their country or with which they have no direct relationship. The CoC tries to address this problem, and its wider adoption should be accelerated.
The representatives of the ODIN, the EC-funded project which uniquely identifies scientists and data sets and connect this information across multiple services noted they managed to roll out a service that requires attributes from the participating entities and they have no problems in getting them. A concrete step would be to exchange experience with the ODIN group on how they address the attribute release aspect.

Level of Assurance is a challenging issue; the first step should be to understand the exact requirements and the costs associated to support them, rather than simply looking at frameworks. There is a need for clearly documenting processes followed by federations and

IdPs. It was agreed to use REFEDS to shape the space to look at. Some more concrete work should be funded via the EINFRA-7 call, which aims to bring together different e-Infrastructures and user communities.
However due to the complexity of this work and due to the fact that many different stakeholders (IdP, resource providers, e-Infrastructures) should endorse any LoA framework agreed upon, it is expected the LoA work to be a longer term action.

The **collaboration** break out session covered different topics:

- there was consensus that a wide solution for guest IdPs should be provided. Commercial solutions should be considered as well as NRENs-delivered solutions
- End-user experience could be improved; in many cases users get error messages that are incomprehensive. The REFEDS discovery guide was mentioned as an example of a model to follow to address this topic.
- The pros and cons on whether big collaborations should create a federation on their own and then interfederate with eduGAIN were discussed. There was no real agreement, as creating and operating a federation requires skilled resources.
- Federation as a service can provide the technical support for operating a federations, however successful local integration of the technologies requires manpower on the local side. This is an essential feature of federated identity management, where policies and processes should be carried out as close to the users as possible, in line with local regulations and in local language.

These inputs will be taken into consideration preparing the content for GN4 proposal and for the answer for the EINFRA-7 call.

Researchers use several tools that do not have a web interface. Research is ongoing to enable **federated access for non-web** application. The discussion focused on how to cross-reference the technologies, compare them, and consider how to bridge them.

Annex I – e-Infrastructures and the AAIs they offer

Identity federations keep growing, both in terms of institutions (IdPs) and services (SPs) joining national federations and in terms of new identity federations. The number of research and education federations that are now participating in eduGAIN, the inter-federation infrastructure operated by GÉANT, has also grown with a total of 24.

Brook Schofield reported that eduGAIN offers a global authentication infrastructure by interconnecting participating identity federations; SAML2 is the adopted standard. eduGAIN policy framework covers various aspects; one important element is the GÉANT Data Protection Code of Conduct (CoC), that describes how services participating in eduGAIN handle the personal information in attributes they receive from the IdPs.

The next challenge is about expanding the CoC to work outside EU countries. This work is currently being done within the f project carried out by the eduGAIN team in consultation with the REFEDS. Results are expected by the end of 2014, whilst the deployment is planned in 2015-2016.

Niels van Dijk demonstrated OpenConext, the collaborative platform to advance work with group management. OpenConext offers a way to manage groups and hence authorisation on the fly. OpenConext is the underlying platform for the Dutch federation; it is also being used by AARNET (Australia) and JISC (UK). Work to standardise group management and to integrate technical solutions within existing e-Infrastructures is still area of research.

Questions on how licences are handled via OpenConext were asked. OpenConext offers different ways to handle licenses ranking from using institutional credentials to more complex ones. SURFnet is working with the SURFmarket procurement group on standardizing the expression of licensing in attributes to ease its consumption. This aspect is of general interest for the R&E sector and should be addressed at European scale.

Vicente Andreu Navarro gave an overview on STORK2, the European infrastructure to share electronic identities (e-IDs) among several Member States. STORK2 is a three year funded project, which started in 2012. The pilot offers a federated system that includes services (SPs), Identity Providers (IdPs) and attribute providers (APs). Services offered via the pilot include, eLearning and Academic Qualifications; eBanking; Public Services for Business; eHealth.

STORK2 is using consent to transfer user's personal data abroad: no data is sent abroad unless user allows the administration to do so. Consent may be given either implicitly, explicitly for data types, or explicitly with data values.
One of the big challenges of STORK is about handling "business attributes", attributes managed by a certain business sector, which often have a meaning limited to this sector. They are retrieved from various Attribute Providers, which introduces the multi-source and multi-country attribute collection challenge. STORK is trying to define ways to measure the quality of the attribute providers to ensure interoperability.
A discussion followed concerning the interoperability between eduGAIN and STORK. There have been already some tests made between STORK and GÉANT (initially led by a team in RedIRIS, the Spanish NREN). STORK services are reluctant to accept credentials that are not handled via the STORK partners; the focus of the tests is on using STORK's credentials to access services

available via eduGAIN. However, because STORK implementation is based on a particular SAML2 profile, full interoperability between eduGAIN and STORK is yet to be demonstrated.

Daan Broeder, reported on behalf of EUDAT. EUDAT aims to develop an infrastructure to support the data-management lifecycle and preservation. The initial approach followed by EUDAT was to create an infrastructure to satisfy as many requirements as possible for the participating communities and to impose no burden on these communities. This meant supporting different authentication protocols to access EUDAT facilities. A credential conversion system was needed to address this requirement. The initial test with Contrail did not lead to the desired results; the new approach is based on Unity, the Cloud Identity and Federation Management part of the UNICORE grid middleware stack. EUDAT is also considering a more pragmatic approach, based on the support for SAML (federated) identities and on requiring eduPersonPrincipleName (the equivalent of username) and to make use of a guest IdP as an alternative.

Peter Solagna reported on EGI, the European Grid Infrastructure. Most of the Grid services must be accessed via X.509 certificates, which are not very user friendly. Each user is assigned to one or more groups, also called a Virtual Organisations (VOs), which defines the roles of the users for authorisation purposes.
To provide a better user experience, there are mechanisms in place to bridge federated credentials and X.509 certificates (i.e. certificates can be generated on the fly after the user has successfully authenticated to their institution). EGI hopes to strengthen the collaboration with the R&E federations. For this to happen a better coverage of federations to also encompass research and data/storage centres would help. To ensure that authorisation happens based on the information provided by the VOs (rather than only on the attributes provided by the IdPs), distributed attribute authorities that are also connected to users' IdP would be needed.

Annex II – Workshop Agenda

10:30 - 11:10 Introduction

- Introduction by the EC on the aim of the workshop - Carlos Morais-Pires (EC - DG Connect)
- Review of the AAA study: what has happened since and what not and why - Opportunities to work on some of the recommendations in the Horizon2020 calls? Jean-Luc Dorel (EC - DG Connect)
- Setting the AAI Scene
  - Federated access (how it works, IdPs and SPs) and international collaboration - Licia Florio (TERENA)

11:10 - 11:30 BREAK

- Setting the AAI Scene
  - Overview on STORK - Vicente Andreu (Universitat Jaume I)
  - EUDAT - Daan Broeder (CLARIN)
  - EGI - Peter Solagna (EGI)
  - Successful inter-federation showcases and the use-cases they support: OpenConext, eduGAIN - Niels van Dijk (SURFnet) and Brook Schofield (TERENA)

13:00 - 14:00 LUNCH

- e-Researchers requirements:
  - DARIAH - Peter Gietz (DAASI)
  - IGTF/FIM4R - David Groep (Nikhef)
  - CLARIN and DASISH - Dieter Van Uytvanck and Daan Broeder (MPI)
  - ELIXIR - Mikael Linden (CSC)
  - CERN - David Groep on behalf of Romain Wartel (CERN)
  - ESA - Andrea Baldi (ESA)

15:00 - 15:15 BREAK (audience splits into groups while they grab a coffee)

· Roundtable discussion on the following topics:

- drivers & barriers for a cross-sector European federated AAI

- AAI user requirements and future use-cases to be supported by an AAI infrastructure

· Closing remarks by the European Commission

Annex III – List of Attendees

Anand Patil - Dante
Andrea Baldi - ESA
Ann Harding - SWITCH
Antonios Barbas - EC / DG Connect
Brook Schofield - TERENA
Christos Kanellopoulos - Greek Research and Technology Network
Daan Broeder - EUDAT/DASISH/CLARIN
David Groep - Nikhef
Dieter Van Uytvanck - CLARIN ERIC
Frank Vercoulen - Eindhoven University of Technology (NL)
Gergely Sipos - EGI.eu, European Grid Infrastructure
Heather Flanagan - Internet Society
Ilse Koning - SURFnet bv
István Tétényi - MTA SZTAKI
Jarkko Siren - EC
Jean-Luc Dorel - EC
John Chapman - Janet
Josh Brown - ODIN
Kitty Fehringer - European Commission
Laura Rueda - ODIN Project
Laurence Field - CERN
Licia Florio - TERENA
Marco Fargetta - Istituto Nazionale di Fisica Nucleare
Merete Badger - Technical University of Denmark
Michal Vymazal - CESNET, a. l. e.
Mikael Linden - CSC - the Finnish IT Center for Science
Miroslav Milinović - University Computing centre, University of Zagreb (Srce)
Nicole Harris - TERENA
Niels van Dijk - SURFnet bv
Pascal Panneels - Belnet
Peter Gietz - DAASI International/DARIAH
Peter Solagna - EGI.eu
Rainer Hörbe - Identinetics
Susan Reilly - LIBER
Vicente Andreu - Universitat Jaume I
Victoriano Giralt - University of Malaga