7th FIM4R Workshop

Summary of the 7th FIM4R with focus on "moving from pilot to production"
European Space Agency – ESRIN - Frascati 23-24 April 2014

Andrea Baldi, Marco Leonardi, ESA
Bob Jones, CERN
Created 12.05.2014
Last Update 16.05.2014 AB, ML, BJ

Workshop Summary

The 7th Workshop on Federated Identity Management for Research Collaborations (FIM4R) was held in Frascati on 23th-24th April 2014 and hosted by ESA at the ESRIN establishment.
All the material presented at the workshop is available online.

The main objectives of the 7th FIM4R workshop, in addition to the usual goal of bringing together the research communities to discuss about federated identity management, were:

- To get a summary about the recent EU workshop organized in Brussels by the European Commission (in cooperation with of TERENA and GÉANT) to discuss the vision for a global European AAI in the context of the Horizon 2020 and GN4 preparation.
- To report and review the results achieved by the various community's FIM pilots since the 6th Workshop
- To discuss about the necessary steps for the community in order to move from pilots to production and how we intend to work with service providers
- To plan the next phase of the FIM4R work in relation with the FIM RDA interest group established in Dublin in March.

Brussels AAI Workshop

The European Commission Information Society and Media DG organized and hosted an AAI workshop relevant for the preparation of GN4 and Horizon 2020 (where FIM is a topic for the program).

The objectives were to gather different communities together for getting an understanding of the status of the AAI initiatives, to review the recommendations made in the TERENA AAA study, to exploit the opportunity of consulting more stakeholders, to see what they are doing and principally to understand the gaps to be filled in order to prepare the next calls for GN4 and Horizon 2020.

The EC took seriously the recommendation coming from the TERENA study, including guidelines for federation policy. FIM4R, with its strong constituency in the European Research, can influence the commission.

Users submitted their requirements via a Google template and even though the result was not a

surprise, it will be useful for the Commission.

It is interesting to note that:

- It is confirmed that the Federation keeps growing; the EDUGAIN coverage is not still optimal and not all the members have signed federation agreements.
- There are still IDPs that are not part of the federation thus preventing access to services for users.
- Non Web applications and Cloud are still rather complex, despite the technology is already identified (ECP, Moonshot, oAuth2) it is not yet used. Nevertheless Cloud/Storage AAI shall be addressed as a priority.
- On the attributes side: IDPs do not release attributes and the issue linked to the crossing of national borders for attributes is not yet addressed at global level (CoC addresses EU only).
- The Level of Assurance is still an issue: what to do for improving it? Should an assurance profile be defined and ask to IDPs to state their compliance? Should IDPs use a stronger authentication mechanism in order to raise the confidence the SPs have in respect to the IDPs?
- Security is still in the infancy for FIM (See next section for more details about operational security):
  - How do SPs know if IDPs are compromised?
  - There's a lack of traceability of the incidents' responses.
  - IDP operational procedures are missing.
- From the User Survey has emerged that:
  - SPs need better support for FIM integration
  - Coordination among different federations would help
  - Missing support for linking groups and accounts

On the programmatic side it is clear that the Commission is very keen to achieve a better cross-section collaboration and the work shall proceed in the 3 main areas: Policies, Services and Cross Sectors (R&D and Commercial).

Terena, which is leading the preparation of proposals for H2020 (EINFRA call) and GN4, considers the FIM4R forum as an important source of requirements, for this reason it has proposed to have all FIM4R organizations providing the list of the key commercial Service Providers (including cloud services) they would like to see integrated within eduGAIN.

Operational Security

Operational Security is another important key topic that was presented and intensively discussed.

The community needs to address Operational security to assure a safe transition from pilots to production so that the Federated services can be trusted by the users (at this stage there are no operational procedures to address the FIM security; what if one IDP within the federation is compromised or a malicious SPs is introduced in the Federation?).

Cybercrime is today very well organized and malicious services are offered on the market.

With Federations the risks are increased because one single incident can affect multiple administrative domains and there is the necessity of collaboration between different organizations. This is not dramatically different from the existing CERT/CSIRT networks, but the incident response for such scenario is not in place and processes are not yet identified. Sooner or later IDPs and SPs will be impacted and we have to deal with them efficiently. It was perceived that some of the

communities are not mature enough to address the operational security autonomously.
In order to operate an inter-federation it is important to have strong operational collaborations with a quick reaction time.
It was recognized that would be very difficult to share a common policy due to the differences between organizations. How could we converge then?
SCI started some time ago to work on a collaborating infrastructure; It considers essential to build the trust, to develop policies' standards and to have incident response procedures in place in order to handle incidents with the cooperation of all the parties involved within the federation. NREN CSIRTs have collaborated for years and they have the right experience to collaborate to develop new FIM operational security services.

From the discussion emerged that some of the communities are quite far from implementing a reasonable level of security; in any case we know that there is the maturity to recognize that security shall be seriously considered in FIM.
Some of the data providers are providing critical data to the user and security is a key issue for them.
It was pointed out that EDUGAIN can only mandate rules for federation but not for security. Nevertheless there was consensus to define the minimum set of requirement (guidelines) for IDPs/SPs that can be recommended.
This could be an interesting topic for Horizon2020, not to replace the work already planned but to fill the existing gap.

Participants appreciated the security operation guidelines and all of them believe that the community would require something concrete in place for the transition from pilots to production.
Imposing requirements on IDPs is a real challenge, but a minimum set of requirements can be identified and agreed so that there might be better chance to get IDPs in the federation adhering to the guidelines. There is a window of opportunity with the new projects coming up and this subject can be a deliverable of one of this projects.
It was proposed to have a small group of people (involving EDUGAIN and FIM4R staff) writing such a minimum set of requirements starting from the existing CSI document, in order to have the document ready by the next Terena Conference.

Discussion on the FIM4R Vision Statement & Priorities

Discussions continue about the FIM4R vision and pilot projects' status.

It was agreed that the FIM4R statement is still valid after 2 years.

> *A common policy and trust framework for Identity Management based on existing structures and federations either presently in use by or available to the communities.*
> *This framework must provide researchers with unique electronic identities authenticated in multiple administrative domains and across national boundaries that can be used together with community defined attributes to authorize access to digital resources.*

In addition it could be extended to include the following points:

- Lifetime of unique electronic identities to cover whole career of a researcher
- Common policy and trust framework shall include operations

- Authorize access to digital resources may imply legal constraints

Priorities and revision of pilot achievements were also made during the workshop and summarized in the following tables. More details about the pilot progress are collected in Appendix A.

User friendliness (**high**) Support for citizen scientists and researchers without formal association to research labs or university Not yet ☹
Browser & non-browser federated access (**high**) Testing in Pilots
Bridging communities (medium) Bridging is a central issue with an efficient mapping of the respective attributes Not tested in Pilots
Multiple technologies with translators including dynamic issue of credentials (medium) Testing in Pilots
Implementations based on open standard and sustainable with compatible licenses (**high**) Down to two standards which can interoperate
Different Levels of Assurance with provenance (**high**) Credentials need to include the provenance of the level under which it was issued Testing in Pilots
Authorisation under community and/or facility control (**high**) Testing in Pilots
Well defined semantically harmonised attributes (medium) Some success for subset of eduPerson but believe it is better to aim for consistency within a community
Flexible and scalable IdP attribute release policy (medium) Bi-lateral negotiations between all SPs and all IdPs is not a scalable solution Not Yet ☹
Attributes must be able to cross national borders (**high**) Data protection considerations must allow this to happen. Not Yet ☹
Attribute aggregation for authorisation (medium) Attributes need to be aggregated from different sources of authority including federated IdPs and community-based attribute authorities. Works for Active Directory Federation Services
Privacy and data protection addressed with community-wide individual ids (medium) Community specific solutions

FIM4R and RDA

Continuation about FIM4R activities and meetings was another main topic. The future of FIM4R was discussed in relation to the recent establishment of the Research Data Alliance (RDA) FIM Interest group and the first RDA meeting that took place in Dublin in March.
The RDA meeting was attended by around 25 people representing SPs, federation operators but not end users (2/3 from EU and 1/3 outside EU).

There was the proposal to start-up an RDA working group for FIM. It was agreed that a broad set of

FIM issues need attention. The proposal could concentrate on:

1. Extending existing EU pilots to US participants.
2. Performing an inventory of the existing communities, use cases and technology projects.

An interesting challenge was to make RDA-P4 registration possible by using home institutions (but this was not taken as a commitment by RDA).
The RDA WGs are more focused and their objective is to deliver result within 12 – 18 months. Meetings are taking place every six months alternating EU and Non EU countries. Sessions are very short (90 minutes, as opposed to the Fim4R workshop that is 1.5 days). Cross-meeting discussions could be absorbed well in RDA by collocation.
Merging FIM4R to RDA FIM would allow a wider participation from US colleagues to the meetings by grouping in this way more communities all together. At the same extent it may be restrictive to discuss on selected topics (WG mandate) by losing the "open forum" characteristic of FIM4R workshops (where lot of peer-to-peer discussions and coordination take place).
In order to meet both objectives it was proposed and welcomed to co-locate FIM4R and RDA meetings, whenever the meeting is hosted in Europe.
A first attempt will be made for the next RDA plenary 4th that will take place in Amsterdam 22-24 Sept 2014).

Conclusions and Actions

The 7th FIM4R workshop has been again a very interesting forum where the community has presented ideas, has shown progress in the pilots and has started new interesting discussions on key issues by prospecting a fruitful future for the FIM. The community is ready to perform the necessary steps to move from pilots to production. The road ahead is quite challenging but nevertheless the wish and the motivation are high and some concrete actions are already on the table to prepare the next phase.

Considering the importance of the new H2020 and GN4 proposals from the AAI community, it is essential that FIM4R input is provided to assure that our needs are addressed in that scope. **Action:** Input shall be provided to Terena (Licia Florio) for H2020 AAI & GN4 proposals: Each research community shall provide a short list of key commercial Service Providers (including cloud services) they would like to have integrated in eduGAIN. These lists will be collected by FIM4R, analysed and the results sent to Terena.	End May
There was a strong consensus about: • A sufficient level of operation security is fundamental for inter-federation production service. • The lack of minimal requirements for eduGAIN IdPs/SPs poses unacceptable risks on the community • FIM4R should leverage the current practices based on existing	

efforts and expertise • The SCI work is relevant and could perhaps be extended to incorporate FIM Proposal Considering the previous points the FIM4R jointly decided to propose common operational security requirements for AAI components (e.g. IDP/SPs) **Actions:** Romain/Dave to circulate the latest version of the SCI paper Romain/Dave to compose and propose a draft document. In collaboration with Géant/eduGAIN (Leif Nixon/Leif Johansson) the FIM4R community shall give feedback and eventually endorse document by following the approach used for the original FIM4R paper.	Mid May End June End August
In relation with the RDA working group and with next RDA plenary, it was agreed: • To schedule the next FIM4R meeting in Amsterdam and co-locate it with RDA 4th plenary (22-24 Sept 2014) • To formulate RDA Working Group focused on the extension of FIM4R pilots to USA partners and on the adoption of minimal set of security operations requirements for IdPs	

APPENDIX Notes of FIM4R talk highlights

Day1 – "Internal discussion about Users Communities updates"

Bob Jones, CERN – Introduction to the FIM4R.
Covered by the summary.

Licia Florio, TERENA – Updates on Trust and Identity Plans.
Covered by the summary.

Romain Wartel, CERN – Operational Scurity and Security Policies.
Covered by the summary.

Day2 – "Session for representatives from Users Communities"

Nicolaus Hanowsky, ESA – ESA presentation.
Nick Hanowsky, head of the Earth Observation Ground Segment Department, intruded the open part of the 7th FIM4R workshop with a warm welcome to all participants.
A short overview of ESA, ESRIN with focus on the Earth Observation functions performed on site was given. The importance of FIM was highlighted in relation to the intense collaborative program that ESA has undertaken with others international organizations for the execution of important programs like Copernicus, LTDP, and the Thematic exploitation platforms for EO.

Bob Jones, CERN – Introduction to the FIM4R.
Covered by the summary.

Andrea Baldi and Marco Leonardi, ESA – ESA EO FIM and Space Federation evolution.
ESA has presented its SAML/Shibboleth based AAI infrastructure operating since 2011; the evolution of the main components (IDP, LDAP, SPs) for what concerns Auditing, Reporting and IT Redundancy was also described. ESA has presented the ESA EO FIM project plans for short, mid and long terms. The ESA AAI evolutions toward a federating approach will be covered by two main project phases:
Phase 1 with Siemens Romania in charge of studying the setup of an internal ESA federation composed by two different ESA administrative domains. This project will be strategic to pilot the integration of the ESA AAI with the IDEM/eduGAIN federations. Additional pilot will be done in the areas of ECP profile, Attribute Authority and the STS (Security Token Service).
Phase 2 would undertake a proper implementation of some of the pilot result. The concept of FIM Space Federation involving ESA and others international partners organizations (DLR,

EUMETSAT, CNES, etc) was presented as a long term FIM objective.

Maryline Lengert, ESA – Federation of partners: the example of Iceland geohazard Supersite within Helix Nebula.

Helix nebula is a federation of partners to build an infrastructure based on cloud computing serving the science community with a strategic plan endorsed by both the scientific organisations and the industry to setup multi tenant/multi cloud providers and to define policies for trust security and privacy. The consortium has increased a lot since the project start (there are now 14 members). Part of the federation is by industry and another part is provided by public infrastructures like Géant (that is part of the federation). Different pilots are in place with CERN, EMBL and ESA. It will be necessary to federate them using Géant/EDUGAIN. The baseline for authentication is the ESA EO SSO.

Tommi Nyrönen, CSC – ELIXIR.

The AAI in ELIXIR is used for bioinformatics and in particular for applying AAI to dataset analysis needed for personalising medicine, for example matching the treatment to cancer by using omics technologies and data analysis tools on secure clouds. The open but access controlled data stored in ELIXIR nodes is available in the public domain, is very important and some of it will be treated with security and privacy. Elixir AAI is based on a distributed infrastructure, with potentially 3 million users, therefore scalability is an essential feature of the system. AAI has been recognised as a key technical service in the ELIXIR program.
Current plans foresee as next step to have REMS **as a service** from ELIXIR Finland to be used to manage access rights to digital data created from biobanking sample collections and stored in data services suited for sensitive data. Increasing the utilisation and understanding of AAI technology for life science service providers is also a priority (more data access committees to rely on REMS). The blueprint for the final scenario is not yet fully defined. Some of the services are not discipline-specific and REMS is stable and usable also by others communities. Also a cross-disciplinary application of REMS has been planned by FI-CLARIN and FI-CESSDA. There is a great need for authorisation tools allowing SPs to take their decisions and tools for Community Management (groups and VOs).

Jean-Francois Perrin, ILL – UMBRELLA.

Umbrella is a AAI for Photon and Neutron community (FP7 projects contribute to it). It counts 30.000 users moving from one facility to another; users need access to experiments by using a single persistent user id (independent of the institution) to have unified access to all tools associated with the experiments. Umbrella is a bridging software with only one authentication server (this simplifies the support for a single user ID). Services provided to users are: Info & service portal, proposal support, remote experiment access, remote file access and metadata catalogue. There is a strong competition between facilities to get access to resources and therefore it is important that the information received from IDP is restricted to the minimum in order to assure privacy. The authorisation is managed by the facilities. In operation there are:

Basic Umbrella together with Geo DNS (to get access to the national IDP) while in progress there are: a bridging with EDUGAIN to use home institution account, a test with Moonshot via Janet to access remote workstations and an ECP prototype to access the iCAT metadata catalogue.
Next challenge: how to access cloud resources with FIM technology (VM, storage)?
Concerning Moonshot experience it was reported that is technically difficult, it requires client software to be installed and the current software is not mature yet.

Mikael Linden, CSC – CLARIN.

CLARIN (Common Language Resources and Technology Infrastructure) is a large-scale pan-European collaborative environment to coordinate and make language resources and technology available and readily useable for Language and Social Sciences & Humanities researchers. CLARIN service provider's federation (SPF) spread in several European countries. A legal Entity (CLARIC ERIC) establishes contracts with with individual national IDFs. Current status is: 17 SPs signed up with 7 IDF and around 200 relevant IDPs who can connect. CLARIN has its own homeless IDP. Collaboration with DARIAS & CESSDA for sharing resources is in the plan.

Peter Gietz, DAASI – DARIAH.

DARIAH is an ESFRI Digital Research Infrastructure for the Arts and Humanities. The DARIAH Federation is composed of four Virtual Competency Centres: the E-infrastructure VCC, the Research & Education VCC, the Scholarly Content Management VCC and the Advocacy VCC. The AAI is based on LDAP and SAML and makes use of standard shibboleth IDP/SP and attributes queries. The IDP is the attribute authority and groups based authorisation is in place. The current challenges are: eduGAIN has too little outreach, not all the institutions sign federation contracts, many IDPs do not release the necessary attributes. DARIAH is operating with a "homeless IDP" since some time (Homeless users are around 1800). The policy for registration is to have a community email address and/or with identity vetting based on a valid ID card checked by a DARIAH trusted person (registration authority model).
Update of registry is performed by an email sent to see if the user is still there. If the email is returning an error, the user is removed form the registry. DARIAH is configured according to the requirements of CoC. There is a plan to move to eduGAIN announcing that they can login with their home institution account. DARIAH will deploy on DFN AAI infrastructure.

Romain Wartel, CERN – High Energy Physics.

CERN is running with an identity federation based on X.509 certificates since more then 10 years. This was very successful, but has some limitations especially on the user side who has to deal with all the issue associated with certificate management.
CERN needs to be able to enlarge the user community and EDUGAIN was identified as the vehicle for doing this. CERN and WLCG decided to work in collaboration with Switch, the Federation Operator of SWITCHaai. Most of the paper work is done, the announcement will follow soon. The objective is to open some services to the EDUGAIN community under two conditions: 1) signed/opt-in inter-federation access, 2) CERN get assurance about minimum security applied.

Next step is to continue to work with a new web-based pilot involving traditional grid services. Two phases are planned: the first one focused on the connection with EDUGAIN (a use case would be using indigo for conference preparation); the second phase is more complex and will include STS to be implemented (with a lot of technical challenges to be addressed). The back-end based on digital certificate will stay in place, but certificates will be generated after a successful login on federated IDPs. There are of course many open questions concerning policy for using the service, security issues due to the federation and many key technical issues (STS, attributes, trust issues, VO user registrations, multifactor authentications, etc.). A transition shall be planned. It is really an interesting window of opportunities.

Daniela Pohn, Leibniz Supercomputing Center – Géant-TrustBroker.

GÉANT trust broker (GNTB) puts the main focus on Federation from an SP prospective.
It aims at supporting user triggered exchange of metadata, complementing existing Federation services, automation of IDP-SP setup, performing attributes conversions when necessary. Shibboleth will be extended to support the above features via plug-ins.
There are two types of federations national (operated by NREN) and community federation (from project). It is difficult to exchange data between them, either they belong to the same federation or are inter-federated e.g. via EDUGAIN (that is a more complex set-up)
The issue with current approach is that the Inter-federation schema includes the minimum set of common attributes and often SPs need more then what is provided. Administrators need to setup all technical details and this is not simple. GNTB automates most of the technical work and complements EDUGAIN. GNTB might be interesting for smaller community that cannot afford to invest in inter-federation to get automation support for a quick set-up.
Project plans are: end on 2014 – March 2015 GNT3+ preparation, IETF Draft for Summer 2015.
GNTB will be presented at t GTNC2014, more details are available on the GÉANT intranet.

Marcus Hardt, KIT – SAML to LDAP bridging development.

Objectives were to develop a Linux login using SAML for non WEB authentication with an easy and pluggable solution to bring SAML to the commandline.
Use case is a pilot for ssh-login using ECP and federated home-IdPs.
A goal was to use ssh client and server without modification. Instead a new PAM module was developed (in python). This module can handle all, a username/password, a SAML-assertion or an ssh-key. The obtained information is used to verify the users validity (either presenting username/password at the home-IdP or by verifying the presented assertion). Furthermore, group information can be provided, based on the attributes present in the SAML-assertion.

After an initial phase the project moved from PAM to LDAP because of some problems with python on scientific linux. The KIT LDAP FAÇADE was introduced in the design; it behaves like a local LDAP but uses interfaces to SAML IDP/SPs on the backend side to determine a users' validity.

The project has addressed the German problems (Privacy laws & co) for handing off of IDP attributes legally (e.g.: user shall "accept" rather than "agree" on this).

Performing harmonisation between X.509 and SAML, both on authentication and on authorisation, is the next hot topic for the project.

Bob Jones, CERN – Summary of RDA session in Dublin.
Covered by the summary.

List of Participants

Ansari, Salim	ESA	Netherlands
Arezzini, Silvia	INFN	Italy
Baldi, Andrea	ESA	Italy
Bingert, Sven	GWDG	Germany
Dumitru, Edith	Siemens CVC	Romania
Fasanelli, Enrico Maria Vincenzo	I.N.F.N.	Italy
Formisano, Ciro	Engineering Ing. Inf. Spa	Italy
Gianfranceschi, Simone	Intecs SPA	Italy
Gietz, Peter	DAASI International GmbH	Germany
Goor, Erwin	VITO	Belgium
Hämmerle, Lukas	SWITCH/GÉANT	Switzerland
Hardt, Marcus	KIT	Germany
Jones, Bob	CERN	Switzerland
Kelsey, David	STFC - Rutherford Appleton Lab.	United Kingdom
Lengert, Mariline	ESA	Italy
Leonardi, Marco	RHEA	Italy
Linden, Mikael	CSC - IT Center for Science	Finland
Lördal, Anders	SWAMID	Sweden
Manieri, Andrea	Engineering Ing. Inf. Spa	Italy
Nyrönen, Tommi	CSC	Finland
Perrin, Jean Francois	Institut Laue-Langevin	France
Poehn, Daniela	Leibniz Supercomputing Center	Germany
Pohl, Christof	GWDG	Germany
Stannard, Simon	CGI	United Kingdom
Wartel, Romain	CERN	Switzerland