

South Africa's Federated Identity Management Initiative

Siju A. MAMMEN

SANREN, Building 43d, CSIR, Meiring Naude Rd. Brummeria, Pretoria, South Africa
Tel: +27 12 841 4213, Fax +27 12 841 4223, Email: smammen@csir.co.za

Abstract

An overview of the concepts involved with identity federation, focusing specifically on it as a means to complement the existing digital identity management infrastructure for research and educational institutions in South Africa, is provided. The paper juxtaposes traditional identity management ideas to the federated model. A description of the South African initiative to establish a pilot Federation is described. The novelty of the South African initiative is the approach that was followed. Specifically it has been driven in a top down fashion. This is compared with other approaches in establishing Federated Identity Management within the NREN community. The ultimate goal of which is to show the reader that several methods exist to establish a federation.

Keywords

Identity Management, SAML, Federation, Single-sign on

1. Background

Web based services, or services in general, require some form of authentication to provide potential users with access. The classical approach that is followed by most service providers (SPs) is to create and maintain a database of users. This approach has a couple of issues and to investigate these potential shortcomings, the perspectives of the two main parties involved, the users and the service providers, need to be analysed separately. They are briefly discussed below.

1.1 The User's perspective

One of the main issues with each SP having their own database is that the users will have multiple log-on credentials for each of the services that they are subscribed to. This, of course, has security issues since the SPs can implement their user database in a variety of ways, and in general the users cannot be guaranteed that their information is secure. Furthermore, if users resort to reusing their username and passwords for all services, a compromise on any one of these databases could potentially compromise all of their accounts.

While the user takes primary responsibility for ensuring an effective password, the service provider's choice of technologies can affect how easily their systems can be compromised. Ideally the users will prefer to have the most secure system, but sadly not all SPs will be able to cater for such needs.

1.2 The SPs perspective

It is difficult to generalise for all service providers, but at least during the initial stages of getting a service up and running, the focus is more directed to the service itself and not on the user access to the service. User access in this case becomes a necessary add on for most service providers. It becomes possible in such situations for service providers to have less than the most secure user database.

Another aspect to consider for the service provider is that the service provider is primarily interested in the attributes of users. The service provider can benefit greatly from reliable, accurate information of its users. However, in the classical sense, the SP has no means to validate the user's attributes since all of these attributes are "self-asserted." And self-asserted attributes limit the type of services the service provider can offer. For instance, if an airline wanted to provide users with a discount if they are a student, this could not be easily catered for with the classical attribute scenario.

1.3 The Third Perspective

Although the user and the service provider are the two main stakeholders in the classical scenario, one further role-player needs to be considered. This is the identity provider (IdP). Most institutions create and manage their employees' user credentials for everyday information and communication technology (ICT) tasks. Creating a user credential would generally involve the collection of various attributes from the user himself/herself, but with the added requirement that the user's attributes are verified in some way.

A lot of the issues discussed in the classical scenario, where the service itself maintains the users' access credentials, could be overcome if SP allowed users to be able to access their service using their work username and password. In essence, each institute or company becomes a party that could provide identities to services providers, hence the term "identity provider." However, while theoretically using the credentials from an identity provider seems like a simple solution, it comes with its own set of restrictions. These restrictions are mainly legal that come about due to the nature of digital identities in general. Institutions need to ensure that they comply with the various legal requirements around identities and will therefore not be able to simply provide their identities to any service provider that requests it.

In most circumstances, when an institution requires a service from a certain service provider, a one-on-one agreement between the two of them is created and the institution's users will have access to that service alone. Leaving aside the scalability concerns for such bilateral agreements, such a model would automatically exclude smaller services that are only needed by select staff to be integrated to the institution's systems.

2. Federated solution

A Federated solution to identity management is one possible framework to allow the credentials held by identity providers to be used by service providers. In essence it is a trust framework between SPs and IdPs to allow the user credentials from IdPs to be used at SPs. In other words, it is a group of identity and service providers that come together to share identities between one

another, where an agreed to set of policies and technologies govern the interaction between these entities.

The policy documents take into account the legal requirements that the IdPs and SPs need to comply with and thereby allow the IdPs to participate in the federation without any legal barriers. And so, when an IdP joins the federation, they will have access to all services within the federation with the ease of mind that the services will not abuse their user's credentials.

1.1. Advantages

The advantages of a federated environment stem directly from the discussion in the previous section:

For the user, he/she can access multiple services with only the username and password from his home institution. Additionally, he can be assured that all his interactions are as secure as his home institution's systems.

The service providers have access to verified user attributes from identity providers.

The identity providers can provide their users access to almost any service within the federation without any additional overhead for the ICT departments within the organisation. Even small obscure services can easily be accessed in a federated manner without hassle.

1.2. Added Features of a Federated Identity environment

In addition to the advantages discussed above, Identity federations introduce several features into a normal identity management system, which would otherwise be unavailable. The 2 most prominent of such features are single-sign-on and federated provisioning (van Vooren, 2007).

1.2.1. Single-Sign-on

This is arguably one of the biggest potential advantages of a Federated model, at least from a user's perspective. The idea is that, within a single web session, the user will have access to multiple services without having to login to each of them separately. Of course there are various challenges involved as well, especially if a user has multiple accounts that he or she wishes to use for different services. Additionally, since not all the services require the same attributes, the Federation should have clear rules in how to deal with this (van Vooren, 2007).

1.2.2. Federated provisioning

Federated provisioning is the exchange of attributes from an identity provider to a service provider in a federated environment. Even though authentication and authorisation occurs via an identity provider, the service provider, in many cases would want to store attributes of individuals accessing their system. The service provider uses these attributes to create local accounts of user profiles in their underlying system. In a federated environment this exchange can be automatically provisioned from the identity provider to the service provider without the user having to intervene. This greatly simplifies the tasks of having to enter one's attributes every time he/she wants access to a service (van Vooren, 2007).

1.3. Limitations

The main limitation for an identity federation at the moment is that the technology is only mature for web-based services. For non-web based services, development is on-going, and an operational non-web based services being deployed in a federated manner may be realised within the next few years (TERENA, 2012). Practically, the lack of non-web based services is not too much of a drawback for most federations as most services will not have a difficult time in presenting their service through a web front end, although in many cases this would be an inelegant way of solving the problem.

3. Research and Education (R&E) Identity Federations

Generally, the scope of the federation is limited to a certain field, both to allow for a more manageable federation and to better focus its policies around a certain community. The research and higher education community within a country is an example of such a limited field. Globally, the R&E community is second only to the public sector in rolling out federations (Hoerbe, 2012). They have been one of the active communities in contributing to the development of technologies relating to identity federations (van Vooren, 2007). The well-defined nature of the research and education community has probably helped drive this trend. The map below shows the countries with pilot and active federations globally. As can clearly be seen, the continent of Africa currently does not host any R&E federations.



Figure 1: Diagram indicating countries with R&E Identity Federations (adapted from REFEDS 2013)

1.4. NRENs and Federations

It is often the case that a single organisation drives the adoption of federation within a community, and within the research and education sphere that organisation is usually a country's national research and education network (NREN). Due to its mandate, relationship and positioning with a country's research and educational institutions the NREN is ideally suited to drive and host the federation services of a country. There are, however, a number of notable exceptions to this rule such as the Australian Access federation which is its own standalone entity, separate from the NREN of Australia.

4. Implementing an identity Federation

In general, an Identity Federation is formed when a group of institutions decide to come together and establish a circle of trust to access services. This circle of trust depends on a number of underlying technologies to enable it. These underlying technologies, the basic rules governing it as well as some legal aspects of forming it, are documented to create the policies of the federation.

The Federation is almost never static and members can enter, or even leave it. Additionally, the policies can need revision over time and some services needed for the federation, like the discovery service, should not be part of any of the member institutions of the Federation. Given these reasons, a formal Federation operator can be formed to manage the network.

The following is a list of the main action steps required to kick-start an Identity Federation initiative:

- ✧ Identify the scope of the Federation
The scope of the Federation generally identifies the institutions that will become part of the Federation as well as the requirements for future institutions to join. For the research and education field, this is generally quite easy to define.
- ✧ Choose a protocol to use within the Federation
The protocol used is the heart of the Federation and the choice of the protocol will depend on various considerations including security, ease of implementation, performance, possible future inter-federation options, etc. In the research and education sphere, SAML is used exclusively in implemented federations.
- ✧ Identify a schema or a set of attributes to be used within the Federation
As mentioned previously, the attributes need to be agreed to between the SP and IdP. And in general it is ideal to standardise on a set of attributes for the Federation. EduPerson is one of the schemas that are popular in the research and education fields.
- ✧ Decide on the architecture of the Federation (optional)
This is a decision that may be made at the beginning of a federation, or later during its implementation phase. In essence the decision comes down to whether or not a central federation agent is recommended for the federation in question. This will be decided by preference and sometimes by legal requirements. In any case, it is important to decide on this as early as possible especially if a hub-and-spoke architecture is decided on since a federation agent should be formed to manage it.

△ Define the policies for the Federation

While the protocols and schema provide the technical foundation for federation, the policies provide the trust mechanisms within the Federation and therefore ensure that information can be used between institutions. In general the policies identify how information can be used, what can and cannot be done with user's information, what the responsibilities of each party in the federation is, etc. Most importantly though, the policies provide the legal framework to ensure that infringement of the policies can be punished.

Once a group of organisations have defined the above 5 items, the only thing left to do is implement the Federation. The large scale implementation is usually a slow task, and could take several years depending on the enthusiasm of the member institutions. However, the initial stages are the biggest hurdle in these kinds of initiatives and as long as a commitment to form a federation is in place, its implementation can occur in a phased manner.

5. Implementing the South African Federation

Until recently each research and higher education institution in South Africa has been operating autonomously without any real need for a federated infrastructure. However, a few institutions recently started projects that allowed them to share resources between one another. And while this project in itself did not warrant a federation to be created, these institutions saw an opportunity to federate their identity management to allow easy access to these shared resources. This led to the South African R&E Federation initiative to be kicked off.

A workshop was held from the 6th to the 9th May 2013 in Stellenbosch, Western Cape to introduce the research and higher education's ICT top management about federated identity management. The workshop aimed to be an educational event with experts from the international community explaining the intricacies of rolling-out a federation. Overall the workshop was well received and the community showed an overall positive attitude to the initiative.

5.1 Long term objective

One of the main outcomes of the workshop described above was the eventual creation of a South African Identity Federation together with all the governance, policies and technologies well documented and functional. The idea is that a self-sufficient entity will be formed that will be driven by community. This is in keeping with the idea that a Federation is a community initiative that no one entity has control over.

5.2 Pilot project

The first step in the long term goal of building an active federation is to run a pilot project with the institutions that are willing to participate. Since the task of getting a federation up and running is a complex task with various aspects to consider, the project structure shown in Figure 2 was decided on by the community.

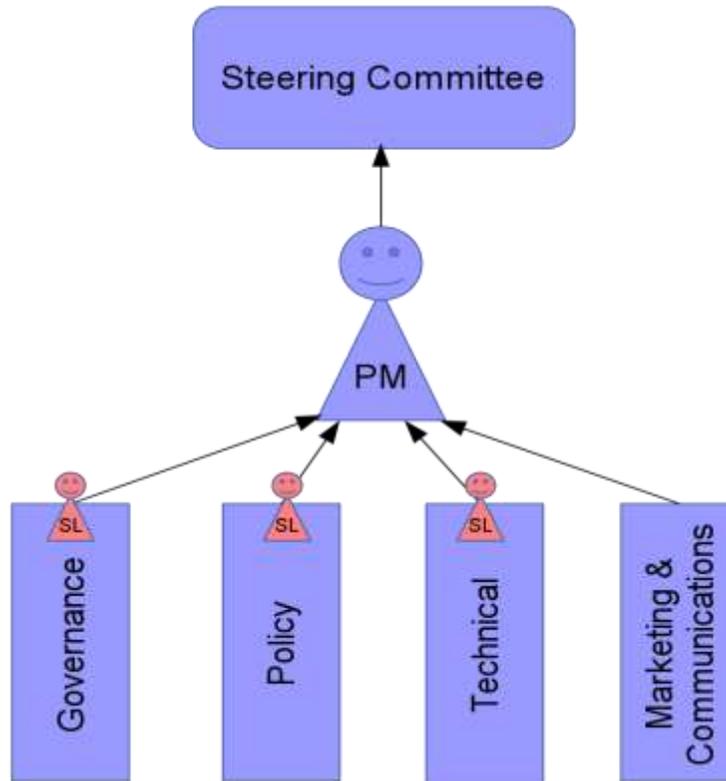


Figure 2: Project governance structure for identity federation project

The outcome of this pilot project will be a test federation together with a proposal from the project steering committee that will be taken to institutions with the goal that interested institutions will join into the production federation.

1.4.1. Roles and Responsibilities

Each of the streams in the project governance structure was tasked with various actions. These are briefly described.

Responsibilities of the Project Manager:

- Managing the various streams
- Reporting progress of the project to the steering committee
- Drawing up a proposal for taking the Federation into production

The Governance stream will define and/or provide recommendations on the following:

- Business Model and Strategy for the final federation
- Model that will be used to fund the production federation
- Scope/Boundaries of the federation

The Policy stream's has been tasked to:

- Draft an initial policy
- Decide on the attributes/attribute release policy needed for the federation

- Define the Level of Assurance (LOA) needed from institutions
- Decide how consent will be handled
- Define the roles and responsibilities of each member institution.
- Decide whether inter-federation needs to be considered from the outset, and what the implications of that decision will be.

The Technical stream has been tasked to familiarise themselves with the available technologies in rolling out Federation and make recommendations on:

- Protocol for the Federation.
- Architecture for the Federation.
- Implementing a pilot of willing institutions.
- Identify potential use cases for the Federation.
- How to inter-federate with other countries?
- The schema used for the federation.

The responsibilities of the Marketing and Communications stream are to:

- Define the Vision/Mission of the Federation
- Draw up official communications to stakeholders
- Help sell the Federation at the right level

The Steering Committee is officially responsible for taking the Federation forward in South Africa. Specifically, they need to:

- Guide the PM and streams to move federation from a conceptual phase through a pilot phase and eventually into a production system.
- Ratify the proposal that can be submitted to member institutions regarding Federation.

5.3 Envisioned challenges

Specifically for South Africa, a number of challenges have been identified in making this federation a success. It should however be noted that these challenges will not be applicable to all potential institutions and it will be an ongoing effort to grow the federation in the future. These challenges are stated here to give an idea of the potential environment that this Federation will be implemented in:

1. **ICT Culture:** As a gross over generalisation, the ICT environment in South Africa (and globally) is conservative and risk averse. It will be a challenge to show that sharing an institution's identities will be beneficial in the long run.
2. **Legal issues:** This is a grey area at the time of that this document is written. South Africa's laws around identity and information privacy are primarily geared to the non-digital sphere. This is slowly changing with the introduction of recent legislation that is said to deal with digital information and privacy more specifically. These laws need to be taken into account when implementing a federation. In general the federation needs to take account of user consent, user privacy, accuracy of identity information, and security of processes.
3. **Funding:** As stated previously, the governance working group shall identify the best model that needs to be employed to fund the federation in the long term. At least for the short term though, someone is needed to fund the initiative. At the moment, it looks like

the South African NREN will fund the project in the short term.

4. Knowledge: While dissemination of knowledge through the introductory workshop and at other conferences has occurred, a ramping up of these efforts is needed as most of the community is in the dark about federation. In particular, it is envisaged that informing the end users about Federation will be challenging. This will be one of the focus areas that the marketing and communications streams will need to overcome.

6. Implementation approaches

The initiative kicked off in South Africa followed a top down approach where the focus was on getting IdPs to be interested in a federated model by getting buy-in from the ICT management within the IdPs. This is not the only approach that can be followed and many existing federations were grown organically from the ground up.

As discussed previously, any SP and IdP can form bilateral agreements with one another to grant an IdP's users access to the SP's service. However, since there are only a few technical options to implement access control between an SP and an IdP, the parties involved would generally opt to use methods that they have had prior experience in. Hence standards such as OAuth and SAML become widely deployed.

The creation of a federation can be easily pictured as the natural evolution of such a scenario where IdPs and SPs, with a narrow focus, decide to formalise the informal decision to stick with a simple standard. This of course is a greatly simplified analogy, the point of which is to make it clear that the bottom up approach can result in the creation of an identity federation as well.

Generally, both the bottom-up approach, and the top down approach is necessary to create a federation. The bottom up approach is usually driven by the technology and the top-down approach focusses largely on the policies and governance, but as we have seen before, both the policy and the technology is necessary to create a working federation.

In most cases where the bottom up approach is followed, the IdPs and SPs that already use a given technology acts as the pilot implementation for the federation. In contrast, when a top down approach is followed, a pilot implementation will need to be created with a group of IdPs and SPs as a proof of concept.

Practically the top down approach is necessary when the institutions' approach to access control for services is widely varied and disjoint, e.g. South African Universities. The bottom up approach comes about, when a widely distributed service requires a standardised form of access control to their systems from multiple institutions, e.g. the Grid. However, there is no "correct" approach and federation is an example of when the path that is chosen is not as important as the destination that is reached.

7. Conclusion

Currently South Africa is well suited to begin the discussion around forming a federation. The country's NREN has been steadily growing the research and higher education network in the country. Additionally, there is sufficient need from the community to form a federated identity

management system, especially from large science projects like the Large Hadron Collider (LHC), Square Kilometer Array (SKA) and South African GRID initiative. And finally, the technologies relating to Identity Federation have become quite mature and well tested and support from the rest of the community is readily available if needed.

The project team created for this initiative has the community support that it requires as well as a sound project governance structure which should contribute to the success of the project.

References

Hoerbe, R. (2012) *Global Trust Framework Survey -DG – Business Cases for Trusted Federations*. [Online] Available from: <http://kantarainitiative.org/confluence/display/bctf/Global+Trust+Framework+Survey> [accessed: 14th July 2013]

REFEDS. (2013) *REFEDS Resources*. [Online] <https://refeds.org/resources.html> [accessed 16th July 2013]

TERENA. (2012) *Advancing Technologies and Federating Communities [Online]*. Amsterdam: . Available from www.terena.org/publications/files/2012-AAA-Study-report-final.pdf. [accessed: 12th July 2013]

van Vooren, T. (2007) *Federated Identity: Enabling the Service Chain*. Nieuwegein: Everett.