www.egi.eu

An 'EGI federated Identity platform' for community-specific Virtual Research Environments

Gergely Sipos

EGI.eu, The Netherlands
gergely.sipos@egi.eu

V1.1 - 25-Sep-2012

In a distributed environment many information service—such as e-mail, library databases, data repositories, portals, grid/cloud computing applications—require users to authenticate themselves. Within a single institute an institutional identity management system can simplify this authentication for the users. Rather than having separate credentials for each system, a user can employ a single digital identity to access all resources to which he/she is entitled within the organisation. Federated identity management extends this approach beyond the institutional level, creating a trusted authority for digital identities across multiple organizations. An identity federation is an arrangement made among multiple organisations to let subscribers use the same identification data (typically institutional username-password pairs) to obtain access to the secured resources of all organisations in the group. Within an identity federation participating institutions share identity attributes based on agreed-upon standards, facilitating authentication from other members of the federation and granting appropriate access to online resources. The organisations that provide 'user verification' within the federation are called 'identity providers' (IdPs in short), the organisations that provide services that verified users can access are the 'service providers' (SPs in short).

Federated identity management is one of the few areas where there is a common interest of the largest, multi-national, European scientific collaborations including the ESFRIs. Integrating identity federations with the EGI production infrastructure would enable these communities to use EGI resources and ultimately would help EGI expand its user base. EGI.eu is therefore would like to enable research communities to build and operate community-specific, portal based Virtual Research Environments connected to the EGI production infrastructure (grid/cloud services) and to their own identity federations. This document outlines a possible architecture of a platform that would simplify the implementation of such portal services and identifies some software components from the EGI community that could be used to implement the platform.

The document builds onto the platform based architecture vision of EGI[1] and actually provides a proposal for extending the collaborative infrastructure platform with a service portfolio that would act as a bridge between identity federations and the EGI Core and Cloud Infrastructure platforms. The new platform – to be called 'EGI Federated Identity Platform' in this document – would exist as a service offered by EGI to research communities, who could optionally use this service when they develop and integrate community specific environments with EGI grid/cloud services.

The EGI Core Infrastructure platform (the grid offering) and the EGI Cloud Infrastructure Platform (federated IaaS offering) are using X509 certificates with VOMS credentials for client authentication

[1] Core infrastructure platform; cloud infrastructure platform; collaborative infrastructure platform. See the details in the EGI Technical Roadmap (EGI-InSPIRE D2.31): https://documents.egi.eu/document/1094

and authorisation purposes[2]. The federated identity platform would simplify the integration of identity federations and community-specific Virtual Research Environments with these two offerings.

Within an identity federation research communities can access services with institutional login names and passwords. While identity federations, for example the NREN based federations, attract a growing user base from various research communities, no single solution exists at the moment to integrate a portal based Virtual Research Environment with an identity federation and with any Virtual Organisation hosted on resources of the National Grid Infrastructures (NGI) from EGI. Various NGIs, regions and communities have different solutions to interface identity federations and EGI Virtual Organisations via web portals.

The INFN Catania team from the Italian NGI recently setup the 'Grid Identity Pool' identity federation (GrIDP in short). GrIDP is an open federation with very lightweight processes for IdPs and SPs to join. GrIDP aims to facilitate cross-institutional, cross-national access to e-infrastructure services. The next few sections use the GrIDP federation as an example to outline how an identity federation could be interfaced with EGI services through the proposed EGI Federated Identity Platform, and to actually outline the services that are encapsulated within the platform itself. The number in the brackets after each listed item refers to steps in Figure 1.

1. The IdPs of GrIDP do not perform identity check when they create a new account for a user, so there is no guarantee that the digital identity issued by an IdP matches with the real-life identity of the user who owns this digital identity. The digital identities of GrIDP could be elevated to higher quality identities by an attribute provider service of the EGI Federated Identity Platform that links personal and/or community specific attributes to user identities. The attribute service would be populated by research collaborations (in the same way how VOMS is populated by VOs), enabling these collaborations to store what they want to store about the people they know. Certain attributes (or the presence of them) for a digital identity can indicate a 'strong identity' for SPs, so SPs that recognise these attributes and can trust these digital IDs more than they trust 'attribute-less' IDs. Besides elevating user identities, user/community specific attributes can be also valuable in making fine grained authorisation decisions inside SPs. (Step 2)
2. The grid and cloud services that EGI resource centres provide require clients to present X509 proxies when they request services. To simplify the integration of service providers with EGI grid and cloud services the Federated Identity platform needs a factory service that can issue X509 proxies for any user ID provided by any IdP. The proxy factory can use user identities issued by IdPs and attributes issued by the attribute provider to generate a proxy. Obviously the proxies need to be recognised by EGI sites, therefore the proxies need to be generated from IGTF accredited CAs, and in some cases (for gLite VOs) need to include VOMS extensions. (Step 3)
3. The platform needs to simplify consistent authorisation across sites and across users for a VRE, i.e. the same level of access needs to be granted at every EGI site for a given user and the same level of access must be granted at every EGI site for similar users (e.g to members of the same project). An authorisation policy provider service is required in the

[2] For details on the various implementations of X509 certificates within EGI please refer to document: Authentication solutions in the European Grid Infrastructure - https://documents.egi.eu/document/1178

platform for this purpose. The attribute provider service needs to be populated collaboratively by the research communities and by the sites who provide resources to them (Step 5).

Figure 1 presents the platform in action. The process starts with the user asking for a service that is accessible through a Web portal based Virtual Research Environment[3]. The VRE delivers the service through a process that consists of six steps. The steps are the following:

1. Obtain the user's identity from his/her institutional IdP (Who is the user?)
2. Obtain attributes about the user (What do we know about the user?)
3. Obtain X509 proxy from factory (Do we need to access resource on the user's behalf? Or a robot proxy will do?)
4. Initiate grid/cloud resource access
5. Obtain authorisation policy from policy provider (What is the user allowed to do given his/her identity and/or attributes and/or proxy?)
6. Deliver service

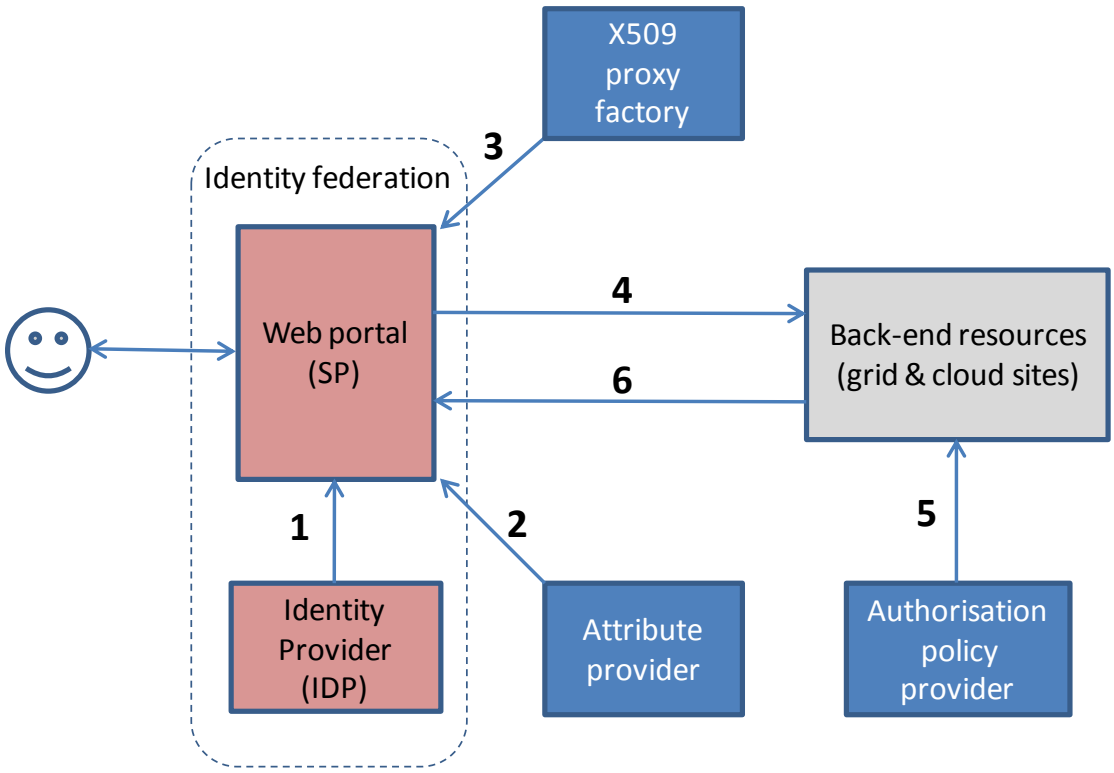

Figure 1. EGI federated identity platform

[3] The document uses a web portal as the a VRE (and therefore the SP in the federation), but it could be as well a service with a programming interface or with a command line interface.

The blue components of Figure 1 represent the services of the platform. The red components belong to the identity federation. The grey component represents the services that EGI resource centres provide and protect with X509 certificates.

After being implemented, the platform will be available for community- specific VREs that need to be integrated with the EGI production infrastructure and need to provide institutional username and password based login. To those who develop these VREs should be able to consume the platform in a 'platform as a service' (PaaS) fashion, with minimal (ideally zero) installation and configuration. Manual configuration of domain specific user registries to act as IdPs in GrIDP, manual configuration of domain specific portals to become SPs is inevitable. Everything else from the platform can be delivered in a PaaS fashion by centrally operated service instances from the NGIs. (Whether one single instance of the Attribute, Authorisation policy and the Proxy factory services could serve the whole community, or multiple instances of these services are needed is not a discussion point until we know the service implementations.)

The platform must be able to provide services for community-specific identity federations, not just to GrIDP. Those communities that already have their own identity federations do not even have a choice and will definitely need the platform to be flexible enough to interact with community-specific identity federations operated by themselves or the NRENs. Services of the platform must be consumable individually, but again without requiring the communities to download, install and operate those services for themselves. In these scenarios the services of the platform should act as independent elements of the 'platform as a service' offering, consumable through e.g. REST APIs.

There are various production level services or service prototypes within EGI that could be used in the platform as the Attribute provider, the X509 proxy factory or as the Authorisation policy provider. One possible platform implementation for example could be:

- Unicore UVOS service hosted at the Polish NGI as the Attribute provider service
- IGI Online CA hosted at the Italian NGI as the X509 proxy factory
- EMI gLite ARGUS service hosted at the Swiss NGI as the Authorisation policy provider

We expect that several presentations of the AAI workshop[4] of EGI Technical Forum will introduce services that can fit into possible implementations of the EGI federated identity platform. Therefore we would like to use the workshop to

1. Refine the above described 'EGI federated identity platform' vision and, if possible, endorse it as a service that the EGI community wants to implement and provide.
2. Identify software providers and service providers from the community that would participate in the implementation and provisioning of a production instance of the platform for the whole community.
3. Identify issues or threats that would make a specific service from the platform, or the platform itself unusable or irrelevant for research communities.

[4] AAI workshop of the EGI Technical Forum 2012: http://go.egi.eu/aaiworkshop