Adopting Identity Federation for the HEP community

Version	Date	Author	Comment
0.1	5th October 2011	Romain Wartel	Initial version
0.2	12th October 2011	Romain Wartel	Incorporated input from Tim Bell and Emmanuel Ormancey
0.3	13th October 2011	Romain Wartel	Incorporated input from Tim Bell and Emmanuel Ormancey
0.4	15th October 2011	Romain Wartel	Incorporated input from Bob Jones and Stefan Lueders
0.5	19th October 2011	Romain Wartel	Incorporated input from Bob Jones and Stefan Lueders
0.6	23th October 2011	Romain Wartel	Incorporated input from Stefan Lueders
0.7	24th October 2011	Romain Wartel	Revised document structure
0.8	25th October 2011	Romain Wartel	Revised document, added input from Dave Kelsey
0.9	28th October 2011	Romain Wartel	Re-scoped document for the identity federation workshop

Table of Content

Introduction

Identity federation is being implemented by a number of organisations and National Research and Education Networks (NRENs) worldwide, as a mean to make authentication and authorization of users across different institutes easier to manage for the resource providers and more convenient for the users. Instead of having one account per resource provider, the users rely on their home institute to act as an Identity Provider (IdP). The resource providers can then handle the user authentication process directly with the IdP, without requiring the user to have a dedicated account for that particular resource provider.

A number of sites would benefit from the gradual adoption of identity federation in the community, in particular:

- External users/visitors would no longer need to have a local account to access resources at a site they are visiting
- Local users would be able to access resources at collaborating institutes using their usual account
- The operational and support costs associated with managing user accounts would be reduced

There are however a number of pending issues and unknowns in order for the community to fully embrace federated identity solutions. This document presents the main challenges towards identity federation in the HEP/academic community, both from a policy and a technical point of view.

Trust Framework

A well-defined framework is required to ensure sufficient trust and security among the different IdPs and relying parties, Anyone can indeed setup an IdP and "offer" to authenticate users from a given site, or *attempt* to pull attributes from that sites, irrespectively of their intent or trust relationship with the site.

The International Grid Trust Federation (IGTF, http://www.igtf.net) is already providing several authentication profiles and some regulation (https://www.eugridpma.org/guidelines) enabling identity providers to become accredited and therefore trusted IdPs. All the Certificates Authorities used in WLCG (100+) are for example accredited. The IGTF has proven to be a reliable, neutral and stable body and provides high quality recommendations and accreditation for authentication in the grid. IGTF-accredited IdPs are for instance required to verify the identity of the end users via a government-issued credential. This elevated Level of Assurance (LoA) is required in order to access grid resources.

The IGTF aims at providing a generic trust framework for all IdPs and considers the fact that most accredited entities are issuing x509-only credentials is simply a technicality. **The IGTF is therefore very well positioned to provide accreditation for non-x509 IdPs**. Although aimed at being technology agnostic, most of the IGTF documents have been written with x509 and public key authentication in mind and **some work is needed to enable the IGTF to support other technologies** like OpenID.

The IGTF should be encouraged to modify the current set of documents, or to produce an additional authentication profile aimed at fully supporting non-x509 IdPs.

The concept of LoA is essential should an organisation decide to open access to its resources to external users. For example, credentials issued by an IdP who verifies the real identity based on a government-issued document and affiliation of its users are more trustworthy than Facebook or Google OpenID credentials. **The level of authorization within the different services at a given organisation should also depend on the LoA of the credentials presented by the user.**

The goal would be for HEP sites to:

- **Enable their IdP(s) to be fully accredited by the IGTF for both x509 and non-x509 authentication systems;**
- **Consider IGTF-accredited IdPs as trusted IdPs, providing an elevated LoA;**

- **Grant a level of trust proportional to the LoA of the credentials presented by the user.**

Attribute management and release

As an identity provider within a federated environment, **each IdP will be asked to release some of these attributes to the external services attempting to authenticate their users**. For example, the remote service may want to check that a given individual really belongs to a particular department or group before granting access to its resources.

Each IdP should specify the attributes that are acceptable to share with third parties, and the attributes that must remain internal to them. Some attributes, for example the real name and the nationality of the user, are particularly sensitive.

Another important aspect of attribute management concerns their verification. In a federated environment, typically, the IdP owns the attributes and release some of them for the resource provider to make an authentication/authorization decision. Authentication and authorization are both handled by querying the IdP. In WLCG, at least one VO has stressed the importance of continuing to verify the attributes of the users autonomously, typically via VOMS (https://twiki.cern.ch/twiki/bin/view/LCG/VOMS). In this model, the IdP would be used for authentication, and the resource provider would contact another entity for authorization.

Security Policies and Procedures

Trustworthy authentication profiles provide an essential framework to issue and manage identities in a federated environment. However, a number of policy issues also need to be addressed to enable the different infrastructure to operate in a federated environment.

In particular, the IdPs in a given infrastructure need to operate in a coherent manner within the boundaries of security policies and procedures addressing the following points:

- Approval of Identity Providers and LoAs
- End user registration and identity vetting
- Roles and responsibility, contact information
- Incident response and vulnerability management
- Service operations
- Logging and traceability
- Attribute release

A revision of the existing policy set in WLCG would need to be conducted to integrate the concept of identity federation.

Furthermore, work is currently in progress (http://www.jspg.org/wiki/Policy_Framework) with several peer-infrastructures (WLCG, EGI, DEISA/PRACE, OSG, TeraGrid/XSEDE) to enable a coherent set of security policies to be adopted by different infrastructures. **Identity Federation clearly suits the unification efforts of the "Security for Collaborating Infrastructures" and should be integrated as part of the on-going work.**

Issues are to be expected with regards to sensitive attribute release. For example, some services require the "nationality" attribute to be presented, while some federations are more conservative and will not even release the real name of the user.

The transition to a federated environment has strong implications for security procedures and incident response in particular. For example, today, each resource provider is responsible to terminate the access of known compromised identities. With

identity federation, **this responsibility will be shifted to the IdP**. The current procedures would need to be revised to ensure the IdPs have the means to provide an adequate response to security incidents and implement the banning of compromised accounts in a reasonable time.
However, the service provider also need be able to retain the ability to revoke any authorization as needed in order to protect their services. Any service provider should be able to block access to any unwelcome user at any time.

Security risks

While identity federation brings a number of advantages, it also introduces a number of security risks that have to be managed. In addition to the implementation vulnerabilities that are to be expected from any new software, identity federation carries several inherent security risks, including:

- Used in the context of the Web browser, when a user is attempting to access a service, identity federation implementations typically prompt the user for credentials either via a sudden redirection to the IdP or via a popup window. In most cases, the user does not type the URL of the IdP. This behaviour creates an **ideal opportunity for phishing attacks**. Phishing attacks are extremely difficult to prevent and are the main cause of compromised accounts at most sites.

- The trust in the federated infrastructure relies on the fact that both service providers and IdPs are trusted. If one of the participants becomes malicious, then other participants might be directly affected. For example, a malicious service provider could pull (and misuse) as many attributes as possible for all the known users at a given IdP. The **implications of operating with malicious participants needs to be evaluated** and understood.

Compatibility and interoperability

The industry is focusing on a relatively small number of technologies to implement identity federation, mostly relying on OpenID, sometimes on Shibboleth. On the other hand, academic implementations are generally based on Shibboleth with support for OpenID in some cases.

It is important to support industry standards wherever possible, in order to increase the interoperability of our services with other infrastructures, while reducing the development and maintenance cost of the adopted solutions.

In addition to closely following the industry standards**, it is important to ensure WLCG and the HEP community have a strategy aligned with other communities**. The more user communities adopt similar technologies, the better chance there is of that they will be supported by commercial software and IdPs.

Technical implications for the infrastructures

Currently, Shibboleth is the main academic solution, while OpenID is the most common solution in the industry.
The industry has focused primarily on Web technologies, with a strong focus on Web browser-based applications.

In distributed infrastructures like WLCG, a high level of trust and fined grained authorization are required, and non-browser-based applications also need to be supported. This indicates that Shibboleth is needed in a number of use cases, while access to traditional or simple services (like Wikis or Web portals) would benefit from the simplicity and ubiquity of OpenID. The community also needs to continue to support x509, which most grid services rely on.

It is not possible to support only one technology, and infrastructures will likely need to manage OpenID, Shibboleth and x509 credentials at least. It may not be possible to concurrently support all technologies for all services however, therefore credential translation services are being implemented to convert the credentials to different formats, for example:

- The EMI Security Token Service (http://www.eu-emi.eu/security)

- Project Moonshot (http://www.project-moonshot.org/)

Such projects aim at providing the missing software components necessary to enable SSO against grid services and convert the credentials to enable interaction between services using different authentication technologies.

In WLCG, users have highlighted the need to hide x509 from end users, who tend to prefer other authentication technologies. In this perspective, converting credentials from a given technology to x509 is required in order to access existing grid services.

While the advantages of such services offering a transparent interface between several industry standards and x509 services are significant for a number of users, there will be operational costs for the underlying infrastructures, which will be proportional to the number of different solutions that need to be supported. Introducing multiple credential translations services within our infrastructure will have operational and support costs that need to be understood and taken into account. Limiting the number of credential translation services and planning for a careful integration within existing services is essential.

An option would be to rely on a central service to issue x509 credentials for grid-specific services, instead of relying on home-issued x509 credentials. IdPs could interface with such a service via a credential translation service. Several services already exist, for example:

- The Terena Certificate Service (http://www.terena.org/activities/tcs/)
- CILogon Service (http://www.cilogon.org/)

The development of credential translation and certificate services should be closely followed, and it is necessary to define the required components needed to enable them to interact with the current x509 services operated in WLCG, at a reasonable operational and support cost.