EGI-InSPIRE

AUTHENTICATION SOLUTIONS IN THE EUROPEAN GRID INFRASTRUCTURE

Date:	**12/07/2012**
Document type:	**REPORT**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	**https://documents.egi.eu/document/1178**

Abstract

This report provides an overview of the various approaches that are currently used within the European Grid Infrastructure to authenticate users. X509 certificates, Terena certificates, limited certificates, robot certificates and identity federation based login mechanisms are introduced and reviewed. The report also provides an analysis of these solutions based on the main criteria that EGI has for an authentication infrastructure before considering it for wider adoption. An action plan that could lead the EGI community to a wide and harmonised adoption of federated identity solutions within the infrastructure is covered by the last part of the report.



I. COPYRIGHT NOTICE

II. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	07/06/2012	First draft	Gergely Sipos / EGI.eu
2	10/07/2012	Minor update based on comments from Roberto Barbera, Nuno Ferreira, Steven Newhouse, Geneviève Romier	Gergely Sipos / EGI.eu

III. APPLICATION AREA

This document is a public report produced by members of the EGI-InSPIRE project.

IV. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE "Document Management Procedure" will be followed:
https://wiki.egi.eu/wiki/Procedures

V. TERMINOLOGY

A complete project glossary is provided at the following page: http://www.egi.eu/about/glossary/.

TABLE OF CONTENTS

1 INTRODUCTION

Resources providers of the European Grid Infrastructure offer services to scientific communities based on the gLite, ARC, Unicore and dCache middleware platforms. Although alternative platforms, primarily the IaaS-like EGI federated cloud platform are emerging, gLite services still dominate the infrastructure (running on more than 90% of the sites). GLite uses proxy certificates for user authentication[1]. Proxy certificates generated from X509 certificates are used as 'tokens' by the job and file management operations performed by the users or by services acting on behalf of the users on grid sites.

During the last decade e-infrastructure communities and their perception of certificate based access has changed significantly. Many of the existing and potential user communities of EGI consider the personal certificate based access as one of the main barriers of uptake. Some of these communities – together with their support teams from the National Grid Infrastructures (NGIs), NRENs and scientific projects – developed various solutions to simplify, sometimes even to completely eliminate certificate based login mechanism for users. Training certificates, Terena certificates, certificate repositories, robot certificates and various types of science gateway frameworks came out from this work. Some of the recent solutions heavily build on 'identity federations' and enable users to access EGI services using their home institutional accounts.

In November 2011 the EGI-InSPIRE project established an 'EGI Virtual Team project' to assess the readiness of the NGIs in adopting federated identity provision mechanisms[2]. The project involved members from five NGIs (Ireland, Czech Republic, France, Switzerland, and Italy) and from EGI.eu. The project was coordinated by a representative of the Czech NGI. The project's scope was assessing the availability of Terena Certificate Service (TCS) and of other federated identity management solutions within the participating NGIs. The work was carried out by defining, then filling out a questionnaire[3] by the participating NGIs[4]. This document summarises the findings of the Virtual Team project survey, and put these into the bigger perspective to define an action plan for EGI towards a harmonised adoption of emerging authentication solutions within the production infrastructure.

Section 2 of the document provides an overview of the various approaches that are currently used within the gLite and ARC middleware platforms of EGI to authenticate users. X509 certificates, Terena certificates, limited certificates, robot certificates and identity federation based login mechanisms are introduced in this section. Section 3 provides an analysis of these solutions. Geographical coverage, science discipline coverage, scalability, robustness, simplicity and integrate-ability with current and emerging EGI platforms are the main criteria for an authentication infrastructure to be considered for adoption within EGI. These aspects are considered in Section 3 for the described solutions.

The aim of this report is to help the EGI community establish an action plan towards a wide and harmonised adoption of federated identity solutions within the infrastructure. Section 4 provides this action plan, which will be further discussed and kicked off at a dedicated workshop of the EGI Technical Forum 2012 event[5].

[1] This is also true for the ARC middleware from the EGI Unified Middleware Distribution.
[2] Federated Identity Providers Assessment EGI Virtual Team: https://wiki.egi.eu/wiki/VT_Federated_Identity_Providers_Assessment
[3] VT questionnaire: https://wiki.egi.eu/wiki/Task_1:_Questionnaire_about_TCS
[4] Answers by the participating NGIs: https://wiki.egi.eu/wiki/VT_Federated_Identity_Providers_Assessment#Actions
[5] EGI AAI workshop: http://go.egi.eu/aaiworkshop

2 IDENTITY MANAGEMENT METHODS IN EGI

This section provides an overview of the various authentication methods that currently exist within the EGI middleware services (gLite, ARC) and provides a brief summary of benefits and disadvantages of each solution. The goal of this section is twofold: First, to serve as a 'white paper' for those who seek for the most suitable authentication method for a service that needs to interact with EGI. Second, to collect the main attributes of the various solutions so those can be further discussed in Section 3.

2.1 Traditional access – personal certificates

The user visits a national/regional Certification Authority (CA)[6] and obtains a personal certificate. The user then joins an EGI Virtual Organisation[7] (VO) that best matches and supports his/her scientific interest[8]. Within the VO the user is identified by the unique name (called Distinguished Name) contained in his/her certificate. He/she can access those EGI sites that allow access to members of the chosen VO.

Pros:

- Users can be personally identified by the grid sites
- CAs that provide personal certificates are available in (almost) every country

Cons:

- Obtaining a certificate is a complicated task for most users
- Obtaining a certificate requires face-to-face meeting with the CA (at one of its registration offices)
- Handling and protecting certificates is difficult

2.2 Terena Certificates

The user requests a personal certificate from the Terena Certificate Service (TCS) provider of his/her country. TCS providers identify the certificate requestor through federated identity mechanism (using the persons' institutional account), they do not require personal visit for identity check. After the certificate is received, the process is the same as for traditional access: user joins an EGI VO that best matches and supports his/her scientific interest. Within the VO the user is identified by a unique name (called Distinguished Name) contained in his/her certificate. He/she can access those EGI sites that allow access to members of the chosen VO.

Pros:

- Users can be personally identified by the grid sites (because a Terena certificate is a certificate that identifies its owner.)
- Obtaining a certificate from a TCS provider is simpler than obtaining it from a traditional CA.
- Obtaining a certificate does not require travel.

Cons:

[6] Certification Authorities recognised by EGI: http://www.igtf.net/
[7] EGI Virtual Organisations: http://operations-portal.egi.eu/vo
[8] A user can be member of multiple VOs at the same time, but his/her work on the grid always happens in the context of one VO at a time.

- Terena Certificate Service providers are not available in every country that provide resources in EGI[9]
- TCS is available only for those who work for an institute that has partnership with the TCS provider of that country[10]
- Handling and protecting certificates is difficult (same as in case of traditional access)

2.3 Limited personal certificates

The user requests a personal, but somehow limited type of certificate from a special CA[11]. Depending on the specialised CA the access to the grid with this certificate is limited in some sense. For example the certificate cannot be used to join any VO, it provides VO membership only for a short period of time, or it can be used only for a limited set of actions within the VO. Typical example of use is for training courses, university courses and service tests. Within the VO the user is identified by a unique name contained in his/her limited personal certificate.

Pros:

- Obtaining a certificate is usually simpler than from traditional CAs
- Obtaining a certificate typically does not require travel

Cons:

- Handling and protecting certificates (files) is difficult
- Certificate is valid only for limited use (VO, time, service)
- Most of the grid sites do not trust these certificates and do not allocate resources for owners of limited certificates

2.4 Robot certificates

Instead of users, the application that these users want to use has a certificate. Users request access to this application and the application accesses EGI sites with its own certificate instead of users' personal certificates. Applications that use robot certificates are typically accessible through a web portal that is already integrated with an EGI VO. The certificate of the application is registered in that VO and the application has access to resources that allow access to members of the VO.

Pros:

- Users do not need personal certificates to access grid resources

Cons:

- Users are not identified individually at the grid level, but they are inside the framework that uses the robot certificate[12]
- Robot certificates are not available in every country
- Cannot be used for applications that accept executables from end users
- Responsibility for user's management is moved to the portal operator.

[9] List of Terena certificate providers: http://www.terena.org/activities/scs/participants.html. Note that some of these providers can issue only 'server certificates' but not 'personal certificates'.

[10] The list of institutes that are eligible to obtain certificates from a given Terena certificate provider can be found on the website of that provider.

[11] There is no up-to-date list of the CAs that provide limited certificates. The GILDA CA is specialised on and issues only limited certificates (https://gilda-security.ct.infn.it/CA/). Some of the national/regional CAs issue limited certificates, but typically in an ad hoc fashion for certain users, groups or events.

[12] The information that must be recorded by the framework (e.g. portal) about the users and their grid-related activities is documented in 'EGI Portal Policy': https://documents.egi.eu/document/80.

2.5 Federated identity based authentication

User has a personal account at his/her home institute (e.g. university) which belongs to an 'identity federation'. The identity federation enables the user to use this institutional account to access services in the federation. There are two main scenarios on how federated authentication can be used in Grids:

(1) Grid sites with their hosted middleware services join identity federations as service providers OR

(2) The Grid middleware services are integrated with the identity federation through intermediary services that translate federated identities to Grid middleware specific identities.

The first option requires significant changes to the middleware and therefore could be achieved only with an enormous development effort. The second case requires much less development effort and can build on top of the existing Grid middleware and operation mechanisms. The identity translation can remain hidden from the user. The federated identity based authentication (either option) has the following benefits and disadvantages:

Pros:

- Users do not need certificates to access grid resources
- Users do not need to apply for additional account to access grid application
- Users can be personally identified at the grid level (depending on how the user's institutional account is mapped to grid certificate)
- Federated model is widely supported outside the Grid community, too.

Cons:

- The notions of identity federations differ slightly among NRENs
- Different identity federations may use different technologies to transfer user account data
- Lack of assessment of the identity providers (similar to how IGTF accredits CAs).
- Identity federations can be connected to current grid middleware services only via identity translator services.

3 ANALYSIS

The 'traditional', X509 personal certificate based access mechanism is available European middleware services for nearly a decade. While it is too technical and complicated from the users' point of view, it satisfies key requirements that EGI has for an identity management framework:

(1) Sustainable: CAs are operated by the NRENs.
(2) Provides wide geographical coverage. CAs are available in Europe, Asia, America.
(3) Provides science discipline-wide coverage: CAs provide services for any user.
(4) Scalable: CAs are established on a per country basis; multiple Registration Authorities can be established for a CA in a large country to reduce travel distances for ID check.
(5) Provides clear methods to report service misuse or abuse: service operators can see who, when and how accessed their sites, they can raise alarms against specific users at the CAs that issued certificates for these users.
(6) Trusted: International Grid Trust Federation provides quality assessment and endorsement of CAs.
(7) It is integrated with EGI middleware services: with gLite and ARC.

While TCS provides a simplified method to request and obtain personal certificates, the fundamentals of certificate management remain unchanged. Users still need to submit the certificate request to a third party: the TCS provider of the country. Users still need to import, export, transform and copy confidential files between browsers, file systems and certificate servers. Terena uses identity federations to simplify the certificate request process. Researchers working for institutes that belong to the federation can quite simply obtain personal certificates from the national TCS provider. Unfortunately this is cold comfort for those who work in a country that has no TCS provider, or work for an institute that is not in the national TCS federation. The Virtual Team showed that TCS is not available in many of the NGIs (see e.g. Ireland, France, Switzerland in Table 1), or for scientific institutes from which NGIs expects users to be affiliated with (see e.g. Italy, Czech Republic). The limited geographical and scientific coverage of TCS are serious limitations for multinational research collaborations, and for EGI too. Because of these limitations TCS can be considered only an extension, but not as an alternative of the CA network. TCS enables simpler access for some of the existing and potential users, but unfortunately it is not available for most of them.

Limited certificates are special type of personal certificates. These are issued by CAs that relax some of the certificate request and distribution rules. For example one can request anonymous certificates (site admins cannot see the user's real identity from these), or the CA distributes a set of certificates to the trainer of a tutorial instead giving these directly to the trainees. Such allowances lower the barrier of infrastructure access, but come with a cost: very few grid sites trust limited certificates and allow VOs with such certificates to use their resources. The sites that yet allow access for such VOs typically do so only for a limited time (e.g. during a tutorial) or to services that can be relatively simply restored without consequences in case of misuse. Because of their nature, limited certificates can be used only for certain use cases in EGI, but not in any generic authentication infrastructure. The GILDA CA[13] provides such certificates in EGI for the whole community.

Robot certificates are personal certificates that are owned by developers of a particular grid application[14]. Instead of using different personal certificates, the application developer's robot certificate is used every time when the application performs a grid operation on behalf of the actual

[13] GILDA CA: http://gilda.ct.infn.it/certification-authority
[14] List of robot certificates used in EGI: https://wiki.egi.eu/wiki/EGI_robot_certificate_users

end users. Because of the single certificate grid sites loose visibility of actual users: they see the load of the application under a single identity. Because of this, robot certificates are allowed to be used for certain types of grid applications: robots can be used only of the application does not accept executable code from end users and runs code that is pre-defined by the application developer[15]. This code is trusted by the application developer and by the CA that issued the robot certificate. Yet, in case of any abuse to grid sites through the application, it is the application developer who has to take responsibility for the security incident. He/she can certainly devolve this responsibility to individual users of his/her application, given that sufficient logging mechanisms are implemented within the application itself so the individual use is recorded at the application level. Although robot certificates have gained popularity within many user communities, they are limited by both their availability and their usability. In terms of availability, at the time of writing, only 10 national/institutional CAs in European provide robot certificates[16]. In terms of usability robot certificates are allowed only for applications that do not take custom executables from their users and access the grid only with pre-defined (and therefore validated) executables.

Federated access to e-infrastructures is the recent and in many respects the most attractive concept for end users. The model could completely eliminate the barriers of e-infrastructure access: a user can use his/her institutional account to connect to services operated by other organisations of the federation. The 'Federated Identity systems for scientific collaborations' workshops[17] and a recent survey run by Terena provided evidences[18] about the fact that scientific communities' prefer federated identity based access with institutional accounts over other means of access. Despite its small size, the VT project well demonstrated the diversity of the EGI community in the uptake of federated identity management solutions. Many of the NGIs and potential EGI user communities do not have access to TCS identity federations (see first, second and third columns of Table 1). Some of the NGIs work on the setup of services that are similar to TCS (see last column of Table 1). Yet another set of NGIs work on bridging technologies to interface national identity federations, or global open identity federations (such as Google, Facebook) to interface portal environments to grid middleware platforms. There is a strong emergence of such bridging solutions within the community, with notable examples provided by INFN-Catania (using Catania Science Gateway Framework), the SCI-BUS project[19] (WS-PGRADE Science Gateway Technology) and the Swiss NGI (GridCertLib[20]). European communities are also active in this area, and run projects that aim to articulate the mutual needs of research and education identity federations worldwide (REFEDS[21]), collect and assess existing AAA (authentication, authorisation and accounting) infrastructures (AAA Study[22]), connect national identity federations into a single international network (eduGAIN[23]), or setup pilot applications that integrate domain services with identity federations (EGA-AAI pilot[24]).

[15] See the details of these use cases in the 'EGI VO Portal Policy': https://documents.egi.eu/document/80
[16] List of CAs that provide robot certificates: https://wiki.egi.eu/wiki/Robot_certificates
[17] https://indico.cern.ch/conferenceDisplay.py?confId=129364, https://indico.cern.ch/conferenceDisplay.py?confId=157486, http://indico.cern.ch/conferenceDisplay.py?confId=177418, http://www.clarin.eu/events/3501
[18] In notes of '37th Terena general assemnbly': http://www.terena.org/about/ga/ga37/CompGA37-6-8.pdf
[19] SCI-BUS project: http://www.sci-bus.eu/
[20] GridCertLib: http://code.google.com/p/gridcertlib/
[21] REFEDS project: http://www.terena.org/activities/refeds/
[22] AAA Study project: https://confluence.terena.org/display/aaastudy/AAA+Study+Home+Page
[23] eduGAIN service: http://www.geant.net/service/edugain/pages/home.aspx
[24] Mentioned in http://dev6.stofnanir.hi.is/is/system/files/IRISC2011-workshop-full-report%20(1).pdf, co-funded from EGI-InSPIRE SA3.

Table 1. Availability of federated identity solutions in VT member NGIs[25]

	Are personal e-science certificates from TCS available in the NGI?	**Are the Grid institutions of the NGI in national TCS federation?**	**Are the institutions of the potential users of your NGI eligible for certificates from TCS?**	**Are there other relevant 'federated identity' based authentication services available in the NGI?**
Ireland	No (but server certificates are)	Yes (for server certificates)	No	Exploring possibilities of a SLCS CA
Czech Rep.	Yes	All major but one (ongoing)	Partly	No
France	No	No	N/A	No
Switzerland	No	All EGI institutions are members	N/A	SLCS (IGTF accredited)
Italy	Yes	Most	Users are expected from outside too.	Preparing a MICS CA

[25] The data is a summary of answers provided by the VT members: https://wiki.egi.eu/wiki/VT_Federated_Identity_Providers_Assessment#Actions

4 NEXT STEPS

Requirements collected from the existing EGI user communities and from potential new user communities clearly show the need for providing federated identity based access to EGI services. The main technologies that enable identity federations in scientific collaborations are developed outside of EGI (e.g. Shibboleth, SAML), and will remain developed outside of EGI. At the same time many of the ecosystem members (NGIs, projects, groups) have technical solutions to interface X.509 certificate based middleware platforms with identity federations. These solutions can be used immediately by national and international VOs to provide federation based login mechanisms for users. A better promotion of these technologies, and community driven further development of these is required. To achieve these goals the EGI.eu Technical Outreach to New Communities team proposes the following action plan:

1. Collect the solutions that exist within the community to interface identity federations with EGI services. Register these solutions in the EGI Applications Database and present them on a new sub-section within the EGI Webpage. (See Science gateways[26] and Workflow[27] as examples for such technical sections.)
2. Facilitate the delivery of training events at EGI Forums, or within NGI events that help the community develop expertise in using, customising the technologies from point 1.
3. Organise topical workshop(s) for the community to discuss
    a. capabilities of existing bridging solutions with respect to emerging needs of scientific communities,
    b. technologies, services and needs that emerge from outside of the EGI community (e.g. from REFEDS, eduGAIN, AAA Study, etc.),
    c. next steps in adopting EGI and external identity federation services within the production infrastructure.

The EGI.eu User Community Support Team and Operations teams will organise a joint topical workshop[28] titled 'Authentication and Authorisation Infrastructure' under the EGI Technical Forum 2012 event (17-21 September 2012) to endorse and kick off this action plan.

[26] EGI Science gateways: http://go.egi.eu/sciencegateways
[27] EGI Workflows: http://go.egi.eu/workflows
[28] EGI AAI workshop: http://go.egi.eu/aaiworkshop