



UNIVERSITÀ DEGLI STUDI DI CATANIA

FACOLTÀ DI SCIENZE MATEMATICHE FISICHE NATURALI

***ENABLING EASY ACCESS
TO GRID STORAGE SERVICES AND DIGITAL REPOSITORIES***

Dott. Antonio Salvatore Calanducci

*“A dissertation submitted in partial fulfillment of the requirements for the degree of “Research
Doctorate in Computer Science”*

Supervisors:

Prof. Giuseppe Pappalardo

Prof. Roberto Barbera

Coordinator:

Prof. Vincenzo Cutello

XXVI CYCLE

THESIS

Enabling easy access to Grid storage services and digital repositories

ABSTRACT

Currently, storage services provided by Grid infrastructures have been mostly used by IT-experts scientists, managers, and researchers due to the their underlying complexities such as protocols, access mechanisms, and user unfriendly interfaces, often based on command line tools. In particular, one of the biggest obstacles that has until now prevented a large uptake of Grid infrastructures by large and diverse worldwide Virtual Research Communities is the Public Key Infrastructure (PKI) on which the Grid Security Infrastructure (GSI) is based. This knowledge is indeed needed for the authentication and the management of X.509 certificates, a mandatory requirement to get access to distributed storage services.

We propose a novel data management architecture, exploiting SAML based authentication, credential delegation, WebDAV and HTTP-standard redirections. This allows users to make use of their institutional credentials to authenticate and provide access to Grid data storage from Web applications, mobile, desktop and built-in clients of popular operating systems. On top of this service, we were able to design a digital asset management system, that promises to simplify the creation, management and access to large scale e-infrastructure based digital repositories aiming at permitting scientists, researchers, students to fully exploit them through tools they use everyday such as Web apps, portals and native mobile apps.

ACKNOWLEDGEMENTS

This work is the result of the many years of study, discussions, research and development carried on at the National Institute of Nuclear Physics in Catania, in the context of many national and worldwide European Commission funded projects. This gave me the chance to travel around the world and to work close to many scientists, researchers and students of different nationalities and culture. This not only had a great impact on my professional carrer but allowed me to grow as human being. So I have to thank INFN for this great and unique opportunity.

I would like to acknowledge Roberto Barbera and my colleagues at INFN in Catania that are my second family and have been able to stand me for so many years in all of my non-sense moments during deadlines and works under pressure. A special thanks to Giuseppe Pappalardo that other than being my supervisor is a friend.

I cannot forget to thank my family that love and support me everyday.

TABLE OF CONTENT

<i>Chapter 1</i>	<i>INTRODUCTION</i>	6
1.1	e-Infrastructures and Grid Middlewares	6
1.2	Grid authentication and PKI.....	7
1.3	Alternatives for Grid authentication	8
1.4	Grid storage implementations.....	9
1.5	Access to data Grid resources with Shibboleth credentials.....	9
1.6	Building digital libraries with Grid storage and metadata services	10
1.7	Structure of this work.....	10
<i>Chapter 2</i>	<i>SIMPLIFY ACCESS TO GRID BY IDENTITY FEDERATIONS AND SCIENCE GATEWAY</i>	12
2.1	Linking Science Gateways to Identity Federations.....	13
2.2	A use case: the DCH-RP e-Culture Science Gateway (eCSG).....	14
2.3	Authentication and Authorization Workflows in Science Gateways.....	16
2.4	Robot certificates and the eToken server.....	18
2.5	User tracking and activity logging.....	21
<i>Chapter 3</i>	<i>GRIDBOX, A RESTFUL DATA MANAGEMENT SERVICE TO ACCESS GRID STORAGE WITH SHIBBOLETH CREDENTIALS</i>	22
3.1	DPM WebDAV implementation.....	23
3.2	Entering GridBox.....	27
3.3	GridBox architecture	27
3.4	Example of GET and PUT operations	32
3.5	GridBox file transfer REST APIs.....	35
3.6	Obtaining Shibboleth sessions tokens for non browser clients.....	37
3.7	WebDAV GridBox interface.....	38
3.8	GridBox file sharing.....	42
3.9	Execution time comparison.....	43
<i>Chapter 4</i>	<i>GLIBRARY, A DIGITAL ASSET MANAGEMENT SYSTEM FOR GRID</i> ...	46
4.1	Features.....	46
4.2	Interaction with gLite/EMI Grid services	51

4.3	Architecture.....	52
4.4	gLibrary first Implementation.....	54
4.5	gLibrary implemetation with GridBox and Federated authentication.....	55
4.6	Web and mobile clients	57
<i>Chapter 5 USE CASES</i>		63
5.1	Data Grids for Cultural Inheritance.....	63
5.1.1	<i>Digitization and data preservation</i>	64
5.1.2	<i>Federico De Roberto works</i>	65
5.1.3	<i>Implementation of the De Roberto digital repository with gLibrary</i>	66
5.2	Data Grids for e-health	69
5.2.1	<i>The DECIDE Science Gateway</i>	73
5.2.2	<i>Authentication and Authorisation of the DECIDE SG with Shibboleth and robot certificates</i>	73
5.2.3	<i>The gLibrary Repository Manager</i>	78
<i>Chapter 6 CONCLUSIONS</i>		85
LIST OF FIGURES		86
BIBLIOGRAPHY.....		88

1.1 e-Infrastructures and Grid Middlewares

In the last 10 years, a new way of doing science is spreading across the world thanks largely to the development of virtual research communities across many geographic and administrative boundaries. A virtual research community is a widely dispersed group of researchers and associated scientific instruments working together in a common virtual environment. This new kind of scientific environment, usually addressed as a “collaboratory”, is based on the availability of high-speed networks and broadband access, advanced virtual tools and Grid-middleware technologies [BBL02] which, altogether, are the elements of the e-Infrastructures [Fos02]. The European Commission has invested heavily in promoting this new way of collaboration among scientists funding several international projects with the aim of creating e-Infrastructures to enable the European Research Area and connect the European researchers with their colleagues based in Africa, Asia and Latin America.

E-infrastructure services are broadly divided in two main areas: computing (also known as Distributed Computing Infrastructures, or DCIs) and storage (also known as Data Grids [CFK+00, SSH+02]).

In particular data Grids provide a set of services that give individuals or groups in a virtual research community the ability to access, modify and transfer an extremely large amount of geographically distributed data for research purposes. Data Grids make this possible through a Grid middleware and services that pull together data and resources from multiple administrative domains and then present it to users upon request. The data in a data Grid can be located at a single site or multiple sites where each site can be its own administrative domain governed by a set of security restrictions as to who may access the data. Likewise, multiple replicas of the data may be distributed throughout the Grid outside of their original administrative domain and the security restrictions placed on the original data for who may access it must be equally applied to the replicas.

Over the years, several Grid middleware stacks have been developed and deployed by resource providers to let their users access the service. After the necessary initial period of research and consolidation of the early middleware stack, a handful of production quality solutions emerged. In Europe, middleware like gLite [LEP+06] from the EGEE project, ARC [EKK+03] from the NorduGrid Collaboration, UNICORE [WS01] and dCache [Fuh04] allowed thousands of scientific researchers to access Grid enabled resources and produce scientific results. Historically the above middleware stacks had been developed simultaneously and even though there was overlap in their capabilities, the delivered solutions were not compatible. Thus usage of these frameworks isolated of the infrastructures and separation of the respective user communities. A clear need for interoperability and standard-based convergence appeared. The growing usage of these software solutions required the transformation of the fragmented European middleware landscape into a harmonized software infrastructure based on professionally managed and standardized services. The European Middleware Initiative (EMI) [AAC+12] was the first project proposed to bring together the four major European Grid middleware providers, ARC, gLite and UNICORE and dCache. This was done in order to capitalize on their long history of competitive development, which has shaped their approaches, but also contributed to the overall quality and clear understanding of key Grid propositions and problems.

1.2 Grid authentication and PKI

All the existing Grid middlewares, and in particular gLite/EMI, rely on the adoption of a Public Key Infrastructure (PKI) [AL99, RSA78] consisting of X.509 compliant digital certificates [CF99] for user authentication and Certification Authorities (CAs) mutually trusted by international Policy Management Authorities (PMAs). All Grid services and interactions need to be authenticated by X.509 certificates or their delegation (proxies). This is a mandatory requirement to authenticate users that require access to DCIs and Data.

Obtaining a certificate requires procedures involving Registration Authorities (RAs) and Certification Authorities (CAs). Moreover managing certificates, with public and private keys, certificate revocation lists and trust keystores, requires knowledge not always available

to non-technical people. This is the first barrier that non-IT scientists and researchers are presented with.

Moreover, working from several locations forces scientists to bring their personal certificate and private key with them. The extra hassle of importing and exporting from one working location to another is impractical even if for Grid and IT experts.

On the other end, the most common and simplest approach to authentication is using the couple of username/password. However, managing credentials could be quite complex and expensive for a distributed environment, especially if resources span resource providers boundaries.

One approach to this problem are federated and user centric identity management systems based on SAML [RHP+08] or Open ID [RR06] that offer Single Sign On and unified Authentication and Authorizaion (AA) mechanism across distributed resources. Shibboleth, a widespread implementation in scientific communities of the SAML standard, allows the user to authenticate with his/her institutional credentials to access Shibboleth protected resources.

However SAML and its implementations, like Shibboleth [MCC+04] , have been designed to be used with the Web, and not with the Grid protocols.

1.3 Alternatives for Grid authentication

Various efforts have been made to integrate a simpler authentication mechanism with Grid infrastructures. One approach is to use a Short Lived Credential Service (SLCS), [Swi13, Tag13] using an online Certification Authority that releases a proxy certificate [WFK+04], a delegated short-lived certificate, to successfully authenticated a user for a short time span. This approach however is easy to use in client environment where proxy certificates are supported, like command line tools or desktop applications. However, most modern browsers are not capable of properly importing and managing proxy certificates. Another slightly different solution developed in the context of the EMI project is the usage of a Security Token Service (STS) [EMI12]. This service implements a SOAP [CLS05] based Web Service, that is able to transform SAML assertions in Grid proxies, using an online CA.

Finally a mechanism designed and implemented by INFN Catania in [LBC+11, BDF+09,], where instead of using online Certification Authorities, the use of Robot certificates [Eur10, CER12] is preferred.

1.4 Grid storage implementations

The EMI project provides Grid data service to store, manage, access and transfer data in the 100 Petabyte range, supporting highly distributed infrastructures. An high-level view of EMI-Data services shows systems both able to keep track on data locations as well as to manage and operate on the associated metadata. The major building blocks, as there are dCache, DPM [ABF+12], StoRM [CFG+07], the LFC [BCL+05], AMGA [KSP08, SK06a, SK06b] and FTS [Mol12], are already in production for several years and with that reached a high degree of stability. One of the main goal of EMI-Data services is to allow costumers to combine those components according to their needs and to build a scalable and easy to maintain data infrastructure.

The Disk Pool Manager (DPM) is a lightweight solution for Grid enabled disk storage management. Operated at more than 240 sites it has the widest distribution of all Grid storage solutions in the EGI infrastructure [KLO10]. It provides an easy way to manage and configure disk pools, and exposes multiple interfaces for data access (rfio [KT01], xroot [DEF+05], nfs [Now89], Gridftp [ABB+03] and http/dav) and control (srm [SSG02]).

1.5 Access to data Grid resources with Shibboleth credentials

The work described in this thesis has focused on how to integrate the Shibboleth authentication mechanism to provide access to Grid data storage. We have designed an architecture where Shibboleth authenticated users could have direct access to the metadata and file systems of Grid storage services without any X.509 certificate or proxy. This allows uploads and downloads from user desktops to Grid resources without caching or streaming data through an intermediate server, strategies adopted by others to solve the same problem.

Our approach makes usage of a credential delegation system based on short-lived tokens together with one of the basic features of the HTTP protocol [FGM+99], link redirection.

The primary advantage of this solution, other than allowing users to access Grid storage resource without any certificate, is the fact that transfers of big files can be handled at the maximum speed allowed by the network link between the user and the destination storage. This avoids the saturation of the bandwidth of an intermediate server that represents a bottleneck in alternative approaches. Moreover, as our solution is based on the HTTP protocol, it has been straightforward to implement a simple WebDAV interface that allows users to move files to and from the Grid, using any WebDAV client, like the ones built into all modern operating systems. Copying a file to a Grid storage is handled with a simple drag and drop operation from a local folder to the mounted WebDAV Grid storage file system.

1.6 Building digital libraries with Grid storage and metadata services

The proposed data management service, named GridBox, forms the basis of a system we have called gLibrary [CCC+07]. This service was designed and implemented over several years to manage distributed repositories of digital objects together with their metadata on e-Infrastructures. gLibrary has been designed with the goal of hiding the complexity of Grid interactions from users. Therefore, no knowledge of certificates or proxy management, storage transfer technologies and protocols is required. New user communities, like the Cultural Heritage, Agronomy, Earth Science, Medical ones, could start using a Grid data services from their browsers and mobile devices to stage and organize collections of manuscripts, plants images, enriched maps with satellite data, PET, MRE, TACs and other medical type images to be shared with their own communities.

1.7 Structure of this work

This work starts introducing the problem of user authentication and authorization on Grid infrastructures. The integration between Federated Identities with Science Gateways, as a solution to the problem, is discussed in the next chapter.

In chapter 3, we describe the architecture and implementation of our novel data management service, GridBox. We compare it to previous works in the same area, demonstrating with some final tests that the overhead of our delegation solution is negligible.

In the following chapter, we present gLibrary, its architecture and implementation

Finally, in the chapter number 5, two real use cases of employing gLibrary in production, one from the Cultural Heritage world and one from the Medical Field are presented.

Chapter 2

SIMPLIFY ACCESS TO GRID BY IDENTITY FEDERATIONS AND SCIENCE GATEWAY

One of the main obstacles for non-IT-expert users to exploit e-Infrastructures, such as Grids, is the fact that they are based on complex security mechanisms such as Public Key Infrastructures (PKI) and generally accessed through low level (command-line based, i.e. non-graphical) user interfaces.

In the recent past, interesting developments have actually been independently carried out by the Grid community with the Science Gateways and by the National Research and Education Networks (NRENs) with the Identity Federations to ease, from one side, the access and use of Grid infrastructures and, from the other side, to increase the number of users authorised to access network-based services.

A Science Gateway is a “*community-developed set of tools, applications, and data that is integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community (US TeraGrid/XSEDE project).*” [WGK+08, Wil07]

An Identity Federation is made of “[...] *the agreements, standards, and technologies that make identity and entitlements portable across autonomous domains (Burton Group)*” [Geb05]. Identity Federations have the aim of setting up and supporting a common framework for different organisations to manage accesses to on-line resources. They are already established in many countries and currently gather a number of people which is in the order of $O(10^7)$.

To make e-Infrastructures easy to use and more accessible, the Italian National Institute of Nuclear Physics is developing since more than two years a Science Gateway

Framework¹ [BFR11] to create a new type of Science Gateways that implements an authentication schema based on Identity Federations.

2.1 Linking Science Gateways to Identity Federations

Identity Federations (IdFs) usually bring together organizations in the field of Education, Research and Culture, namely Universities, Research Institutes, supercomputing centres, medical research centres, National Libraries and Museums, and other cultural institutions.

Organizations subscribing to an IdF link their Identity Provider (IdP) to the Federation. An Identity Provider is a service which enables end users belonging to the organization to use their usual credentials, and more generally their Digital Identity, in order to connect not only to resources provided by their own organization, but also to those offered by other federated organizations. Thanks to the federated approach, once a Science Gateway links to a specific Federation, becoming a Service Provider (SP) of the Federation, all end users belonging to that federation are immediately enabled to be authenticated into the Science Gateway. This does not imply that they are automatically authorised. Indeed, unlike “old fashion”, command line based, access to Grid infrastructure, where X.509 digital certificates and their proxies, possibly containing VOMS [ACC+04] extensions, are used both to authenticate and authorise users, a feature of the INFN Science Gateway Framework is that it decouples the authentication and authorization steps, the first one being demanded to the Federations’s IdPs while the second one remains with resource owners and implements their own access policies. Each user of an IdF will need to be authorised to access a specific resource within a Science Gateway according to its owner’s policies. So, different user groups will access different subsets of resources and may have different rights on them.

¹ The “Catania Science Gatewat Framework”, <http://www.catania-science-gateways.it>

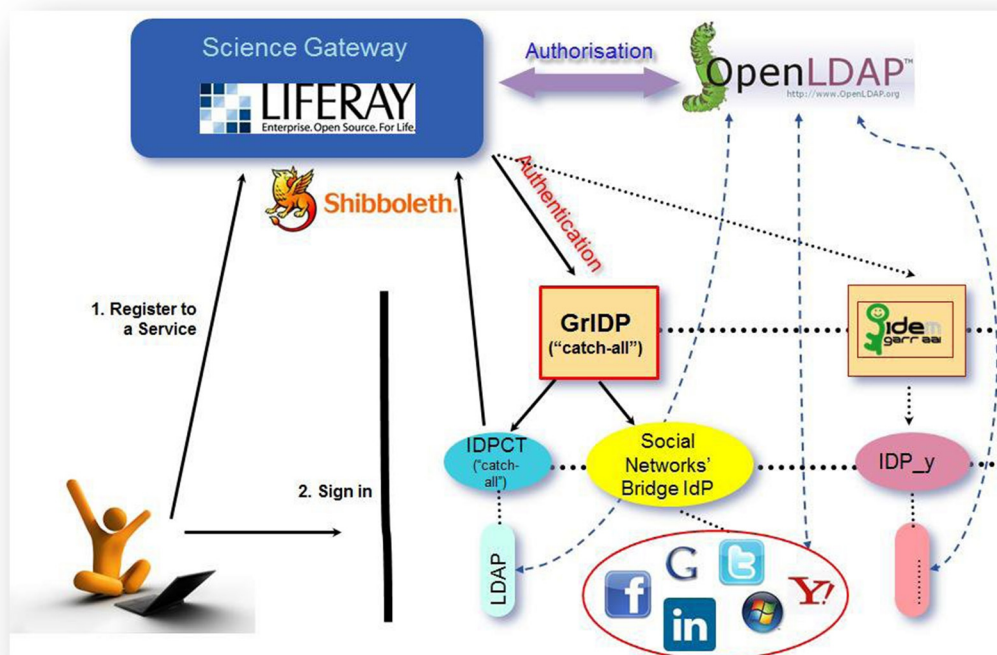


Figure 2. 1 Authentication and Authorization Schema of a Science Gateway

2.2 A use case: the DCH-RP e-Culture Science Gateway (eCSG)

The DCH-RP² Digital Cultural Heritage Roadmap for Preservation is a coordination action supported by EC FP7 e-Infrastructures Programme, launched to look at best practice for preservation standards in use.

The project aim to harmonize data storage and preservation policies in the digital cultural heritage sector; to progress a dialogue and integration among institutions, e-Infrastructures, research and private organisations; to identify models for the governance, maintenance and sustainability of the integrated infrastructure for digital preservation of cultural content.

In the context of the DCH-RP project, an e-Culture Science Gateway (eCSG) has been developed by INFN to demonstrate a model to enable transparent access to Digital Cultural Heritage contents for as many researchers all around the world as possible. To achieve this goal, the e-CSG has been implemented with the INFN Catania Science Gateway Framework that in turns make use of Identity Federations as authentication schema.

² The DCH-RP Project Homepage, <http://www.dch-rp.eu>

So far, the DCH-RP eCSG is integrated in IDEM³, the Italian AAI Federation dedicated to Research, Education and Culture, managed by GARR⁴, and in GridP⁵, a “catch-all” federation also managed by GARR.

In order to access the e-Culture Science Gateway services, a user must be both authenticated and authorized. The schema for authentication and authorization is depicted in Fig. 1. User authentication relies on IdPs that are members of one or more Identity Federations. Currently, the INFN Science Gateway framework only support federations based on the SAML 2.0 standard specifications and on its implementation done by Shibboleth and SimpleSAMLphp.

Other than integrating with the GARR IDEM Federation and, through it, the eduGAIN inter-federation, the eCSG also support all the Identity Providers of the Grid IDentity Pool (GridP), a “catch-all” Identity Federation jointly operated by INFN Catania and GARR that has been expressly created to gather all the IdPs that do not already belong to any official federations and all the users of the eCSG who are not (already) registered in any IdPs. This is particularly important and useful in the contexts where it is necessary to authenticate the so-called “citizen scientists” (i.e., people belonging to the general public) and let him/her access the e-Infrastructure for dissemination and self-learning purposes.

Inside the GridP Federation, two special IdPs have also been created: IdP Open⁶ (a normal Shibboleth-based service) and the “Social Networks’ Bridge Identity Provider”⁷, that allows users to get authenticated with the same credentials they already have with the most known and populated social networks. Both IdPs have recently been endorsed by GARR and are maintained at GARR premises and their availability is in line with the recommendations contained in the recent TERENA AAA Study [EC13], a study on Authentication and Authorisation platforms for scientific resources in Europe, on behalf of the European Commission.

³ *IDEM*: <http://www.idem.garr.it>

⁴ *GARR*: <http://www.garr.it>

⁵ *GridP*: <http://Gridp.garr.it>

⁶ *IdP Open*: <http://www.edugain.org>

⁷ Social Networks’ Bridge Identity Provider, <http://idpsocial.garr.it>

2.3 Authentication and Authorization Workflows in Science Gateways

Unlike authentication, user authorization is carried out at the level of a Science Gateway: users whose request to register is approved by the managers of the SG, are stored in a LDAP-based registry together with the roles they have and the privileges they are granted.

The workflow by which users can register to a Science Gateway is shown in fig. 2.2. A user goes to the Science Gateway and asks to be registered by filling a dedicate web form. Here, she can specify the Identity Federation she belongs or ask to be enrolled as a member of GridP. The request of registration, once it is confirmed by the user via email, is then forwarded to the administrators of the portal. If it is accepted, user information is stored on the LDAP registry and the user is notified that she can sign in. Otherwise, she is notified that her request has been denied. This procedure has been put in place in order to ensure that authorisations are not provided automatically to everybody and that a check be done on the requests by a human being.

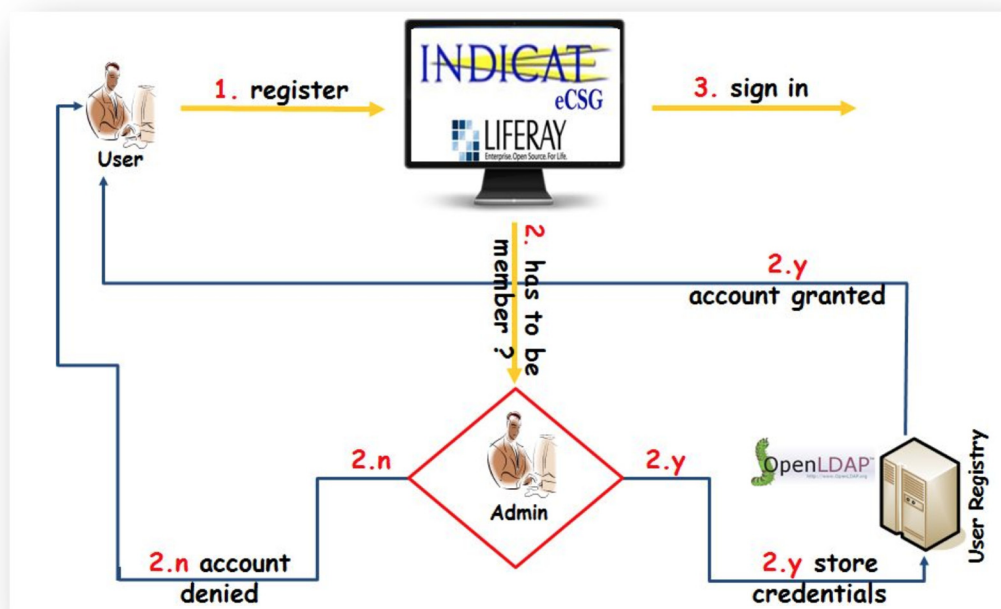


Figure 2. 2 Workflow of the registration procedure

Once a user has been authorised to access the Science Gateway, she can then sign in and run the applications she is allowed to from within the portal. The workflow of this phase is depicted in fig. 2.3. When the user signs in, she is asked to select in a web page the Identity Federation and the Identity Provider she belongs to.

Then, she is redirected to the login page of her Identity Provider where she can insert her credentials. If they are correctly verified, the control returns to the Science Gateway that checks if the user is inserted in the LDAP registry. If she is, the user is then presented with the web page(s) of the application(s) she has the privilege to run on the Grid infrastructure. At this time, the portal contacts an eToken server (more details below) that returns a valid proxy to be used to perform the Grid transaction. Once done, the output of the application is presented back to the user who can thus download it on her computer.

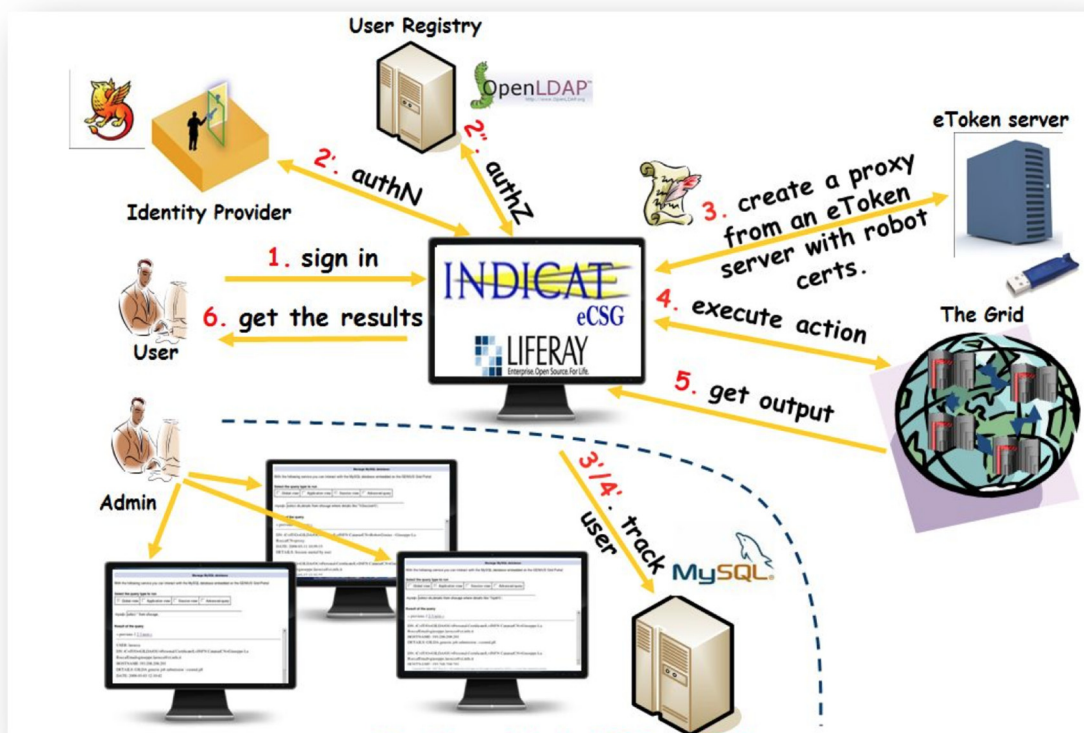


Figure 2. 3 Workflow of the sign-in procedure

2.4 Robot certificates and the eToken server

Obtaining the access to a Science Gateway is only the first part of the authentication process. Users still need to be authenticated by the Grid security mechanism in order to access the available data and computing resources.

The Grid Security Infrastructure (GSI) [FBS06] is based on the adoption of a Public Key Infrastructure (PKI) consisting of X.509 compliant digital certificates for user authentication and Certification Authorities (CAs) mutually trusted by international Policy Management Authorities (PMAs). So users should in principle use their own certificate to access Grid services. Other portal implementations request users to upload proxies or even personal X.509 certificates to the portal, to authorize the following interactions with Grid services accessible through them. This approach is highly not recommended in a Science Gateway environment not only because it usually creates complications for non-experts, but it's considered a very unsecure practice. Private keys, one of the component of a X.509 certificates, should be kept in a safe place a never cross public networks.

A solution to this problem, implemented in the INFN Science Gateway Framework, is to make use of a mechanism that creates proxies on the fly and on user request. This is done by a service called eToken server [LBC+11]. The eToken server generates proxies starting from robot certificates. Robot certificates [Eur10] are special, yet standard, digital certificates stored in USB Smart Card, referred to as e-tokens. Therefore, it is possible to bind robot certificates with applications accessible from a Science Gateway and allow people to run them without any personal credentials. Figure 2.3 and 2.4 shows the workflow and the architecture.

According to the proposed schema, in order to use Grid services, a user has to be authenticated on the Science Gateway, so the log-in depicted in Figure 2.1 comes in action. Then, when an operation to a Grid resource is requested the authorisations of the logged-in user are verified in the LDAP registry and the portal retrieves a valid proxy on behalf of him from the eToken server managing robot certificates.

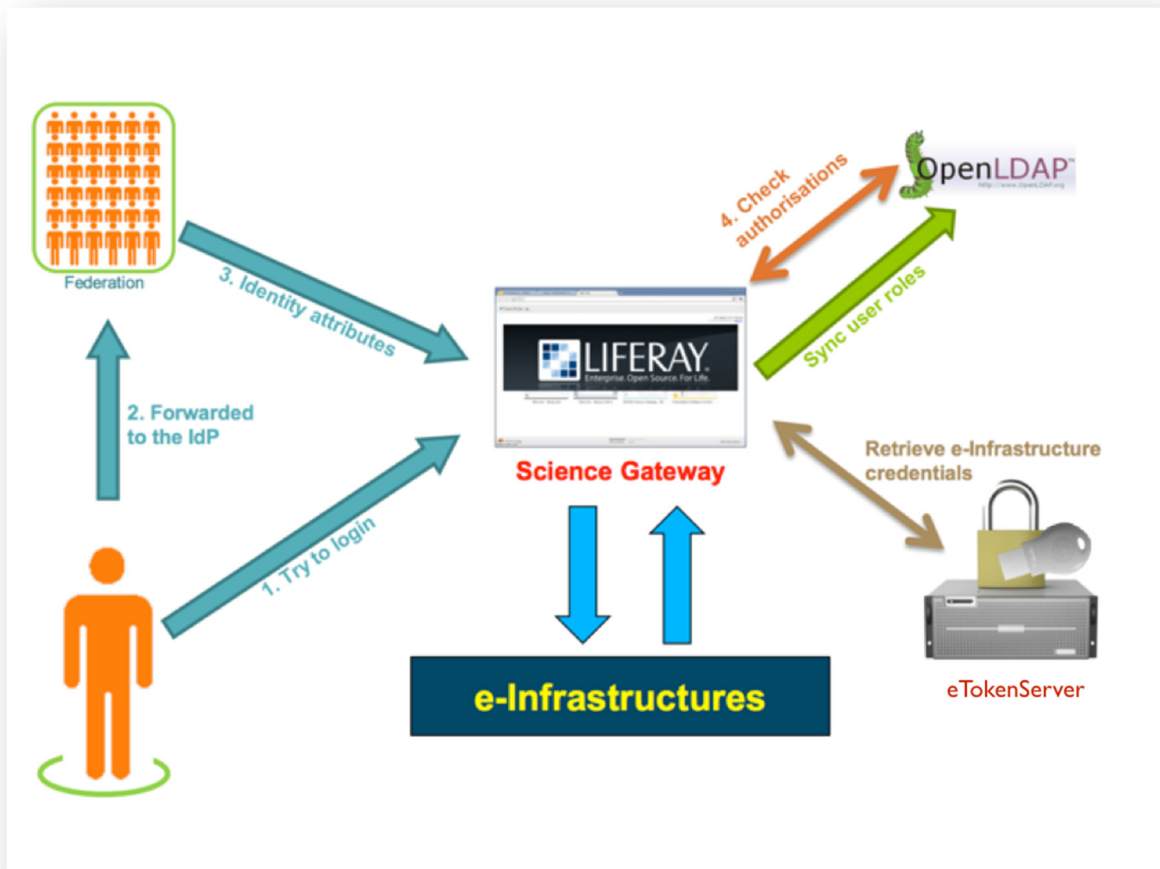


Figure 2. 4 Architecture of the Authentication and Authorization (AA) system in Science Gateways

The proxy generated on the fly contains the extensions that specify the role and privileges of the robot certificate inside the Virtual Organization supported by a given Science Gateway, so different proxies can be created according to the different roles and privileges of the user in the LDAP registry. This ensures a fine grained authorization and provides the portal manager with the complete control of deciding what a given user can see and do.

The core of the eToken server is a “lightweight” Grid crypto library whose Service Oriented Architecture is depicted in fig. 2.5.

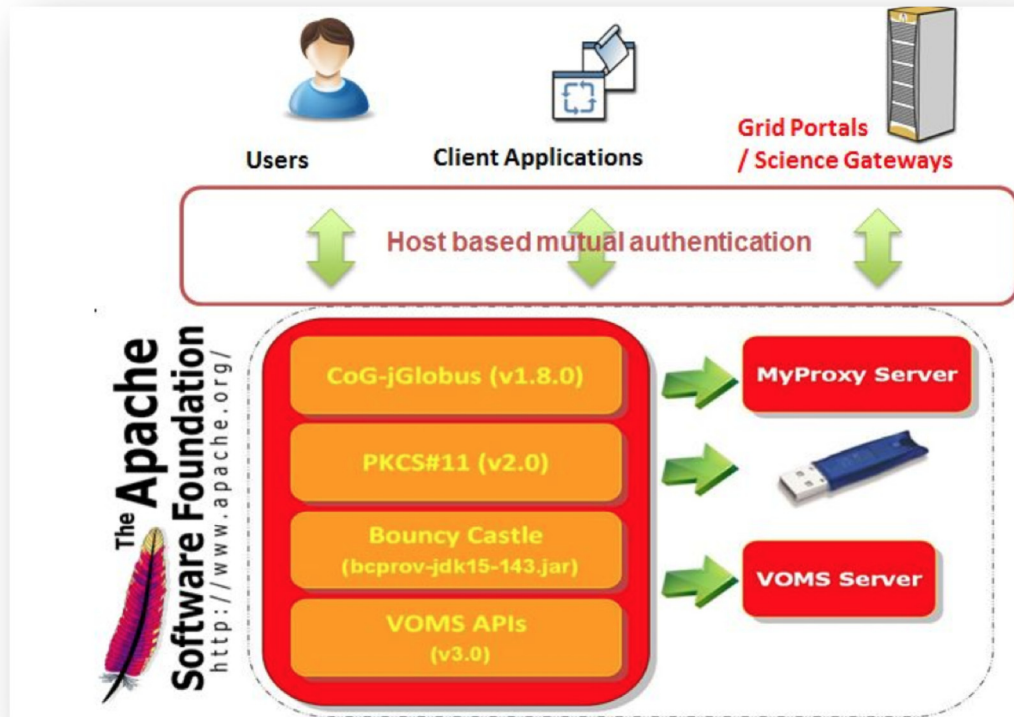


Figure 2. 5 eToken server architecture

The multi-threaded eToken server holds the web services to access the smart cards and interacts both with the Virtual Organisation and the automatic proxy renewal (MyProxy) [BHW05] server. A Java multi-platform client configured for inter-service communication via HTTPS completes the architecture. In order to improve the performances, the server is built on top of the Apache Tomcat⁸ Application Server and configured to accept requests only from a set of authorized “clients” (e.g., the Science Gateway). The adoption of the Apache Tomcat as Application Server ensures scalability and high performances especially when the server has to deal with huge numbers of requests. To further improve its performances and reduce the waiting time to get a proxy, the eToken server implements a mechanism for caching the proxies.

⁸ Apache Tomcat, <http://tomcat.apache.org>

A REST [Fie00] API has also been implemented to access the eToken server. The currently available functions allow to:

- List all the robot certificates stored on the different smart cards physically connected to the server;
- Generate VO-specific proxies from a given robot certificate;
- Create and upload a long-term proxy on a MyProxy Server, allowing proxies to be automatically and securely renewed for much more than their default 12- hours lifetime. Once generated, the proxy signed with the robot certificate is used to entrust the Grid transactions, such as job and data management ones.

2.5 User tracking and activity logging

Actually, proxies used in Grid transactions are not strictly bound to the user since the Subject or Distinguished Name (DN) of the proxy, coming from the robot certificate from which it has been generated, does not contain any information about the user. Hence, from the Grid point of view, a single user (the robot owner) is performing all the different operations made by the portal.

In order to be compliant with the strict rules of the European Grid Infrastructure VO Portal [EGI12] and Grid Security Traceability and Logging [EGI12a] policies, each operation on Grid done on request of the user from the Science Gateway are tracked and stored on a User Tracking System (DB) that can be inspected at any time by the administrator of the portal. This component performs the association between users and Grid operations in a non-repudiable way, combing information coming from the services with the records in the Logging and Bookkeeping service [MKR+07]. This ensures the non-repudiability of Grid transactions, which is one of the most important requirements of the Grid Security Infrastructure.

Chapter 3

GRIDBOX, A RESTFUL DATA MANAGEMENT SERVICE TO ACCESS GRID STORAGE WITH SHIBBOLETH CREDENTIALS

The problem we are trying to solve is to find a mechanism that lets a user access files on Grid storage services without a personal X.509 certificate or proxy, but at the same time have a way to identify users and grant access only to authorized ones.

As discussed in the previous chapter, we have integrated the Shibboleth authentication system with Web portals, based on Liferay, but this could be easily applied to any Web application. Successfully authenticated users are then authorized to the Web portal, contacting a LDAP directory that maps users with proper roles and/or group membership. Users, with roles and groups for which Grid interactions are allowed, will be able to access portal pages and areas (JSR-286 portlets, in the case of INFN Web portals implementations) that requires access to Grid services. These pages and portlets can request the generation of a proxy certificate from the eTokenServer REST APIs. This proxy will be used in the ensuing interactions with the services provided by a Grid infrastructure.

In the case of data management operations such as file downloads, uploads and replication this architecture where a Web portal is the middle tier between the user and Grid storages is unpractical in real scenarios. With this approach, any data movement operation would transit from through the middle tier. This is really inefficient especially for large files transfers. A way to establish a direct connection between the user and the destination storage, without any middle caching was needed.

3.1 DPM WebDAV implementation

In the latest release of DPM, a very popular and widespread Grid storage implementation, a WebDAV interface [ABF+12] has been recently introduced, beside the GSIFTP and RFIO transfer protocols to handle file transfers. However, all the interactions with the storage endpoint are managed over HTTPS, and need to be authenticated by a X.509 certificate or a delegated proxy.

Every DPM deployment is made up of a *head node* and one or more *disk nodes*. The aforementioned WebDAV interface listens for requests on the *head node*, and after the user is successfully authenticated, it redirects the user's client to the disk node that contains the requested file. This request uses either the HTTPS or HTTP protocol to complete. However, there is no authentication using the user certificate to the disk node's HTTPS/HTTP endpoint. The authorization is checked instead by using a short lived token, generated by the *head node*, and appended to a redirection link that is returned by the head node to the client on the first request as showed in the following diagram:

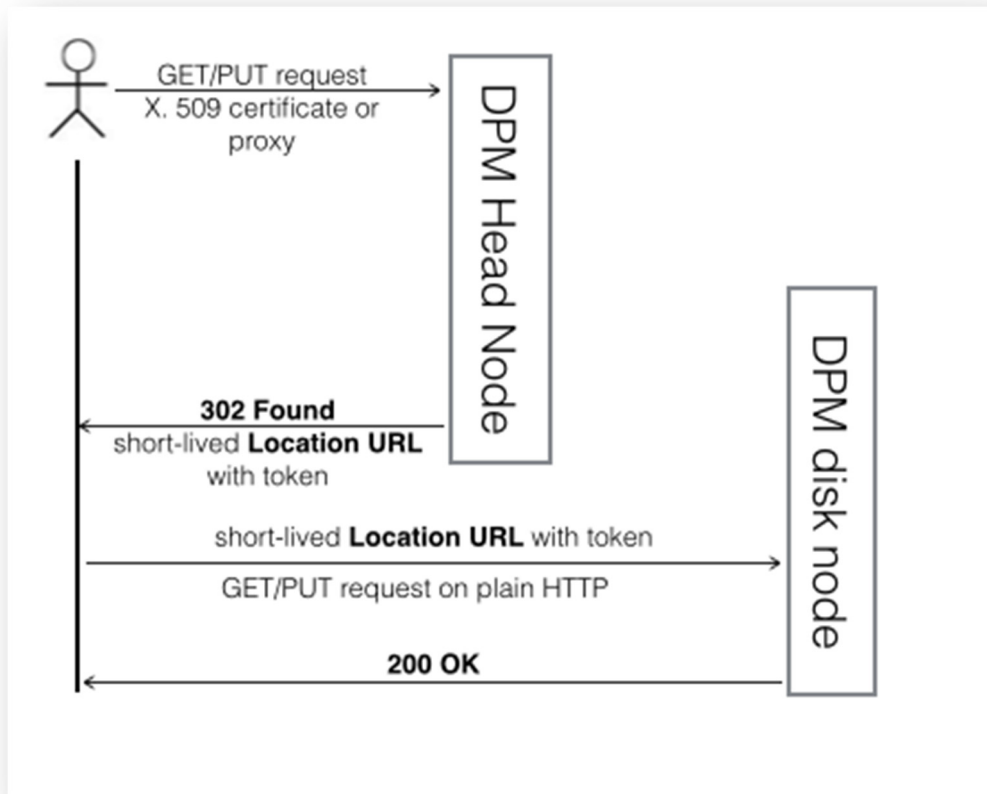


Figure 3. 1 DPM HTTP redirection mechanism

In the following example, we used the command line utility *curl*⁹, an advanced and popular HTTP client, to retrieve a file from a DPM storage element. We are using a proxy certificate for authentication and authorization. As soon as the head node verify the file existence and permissions, it will return a redirect link in the Location header response, that redirects the client to the disk node (in this example both the head node and disk node resides on the same machine):

```
$ curl -O -v -L -E proxy -X GET https://prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-  
rp.eu/test/demo.jpg
```

```
About to connect() to prod-se-03.ct.infn.it port 443 (#0)
```

```
* Server certificate:
```

```
* subject: C=IT; O=INFN; OU=Host; L=Catania; CN=prod-se-03.ct.infn.it
```

⁹ <http://curl.haxx.se>

```

*      start date: 2013-10-02 13:04:30 GMT
*      expire date: 2014-10-02 13:04:30 GMT
*      subjectAltName: prod-se-03.ct.infn.it matched
*      issuer: C=IT; O=INFN; CN=INFN CA
*      SSL certificate verify ok.

> GET /dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg HTTP/1.1
> User-Agent: curl/7.33.0
> Host: prod-se-03.ct.infn.it
> Accept: */*
>
< HTTP/1.1 302 Found
< Date: Fri, 29 Nov 2013 10:53:25 GMT
* Server Apache/2.2.15 (Scientific Linux) is not blacklisted
< Server: Apache/2.2.15 (Scientific Linux)
< Link: <http://prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-
rp.eu/test/demo.jpg?metalink>; rel=describedby; type="application/metalink+xml"
< Location: http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-
29/demo.jpg.23461.0?token=9yStWppfKd3LiJBD2PII6BiNSHk%3D%401385723405%400&dav_s
fn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg
< Vary: Accept-Encoding
< Content-Length: 485
< Content-Type: text/html; charset=iso-8859-1
<
* Connected to prod-se-03.ct.infn.it (193.206.208.163) port 80 (#1)
> GET /storage/vo.dch-rp.eu/2013-11-
29/demo.jpg.23461.0?token=9yStWppfKd3LiJBD2PII6BiNSHk%3D%401385723405%400&dav
_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg HTTP/1.1
> User-Agent: curl/7.33.0
> Host: prod-se-03.ct.infn.it
> Accept: */*
>
< HTTP/1.1 200 OK
< Date: Fri, 29 Nov 2013 10:53:25 GMT
* Server Apache/2.2.15 (Scientific Linux) is not blacklisted
< Server: Apache/2.2.15 (Scientific Linux)
< Content-Length: 8949
< Content-Disposition: filename="demo.jpg"
< Accept-Ranges: bytes
< Access-Control-Allow-Origin: *
< Access-Control-Allow-Methods: POST,GET,DELETE,PUT,OPTIONS,TRACE
< Access-Control-Allow-Headers: Content-Type,Content-Disposition,X-Requested-With,X-File-
Type,X-File-Name,X-File-Size
< Content-Type: image/jpeg
<
{ [data not shown]

```

Table 3.1 File download via HTTP GET using a proxy certificate

This token augmented redirect URL, however, become invalid after a few seconds and the disk node will return a 403 Forbidden status after its expiration. So the requesting client only has a few seconds to follow the redirect. Moreover, even if this redirect link reaches an unauthorized user, the server will return the same error as the link has to be followed only by the client that authenticates with the head node first and initiates the request. This is implemented using the IP address of the requesting client as a key to generate the token. The disk node, decoding this token, will expect that the decoded IP address matches the one of the requesting client.

The following example demonstrates how the direct usage of returned redirect URL (if requested after some seconds and/or from a different IP address) doesn't automatically grant access to the requested file in the DPM disk name space:

```
$ curl -v http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-29/demo.jpg.23461.0?token=9yStWppfKd3LiJBD2PII6BiNSHk%3D%401385723405%400&dav\_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg
* About to connect() to prod-se-03.ct.infn.it port 80
* Trying 193.206.208.163... connected
* Connected to prod-se-03.ct.infn.it (193.206.208.163) port 80
> GET /storage/vo.dch-rp.eu/2013-11-29/demo.jpg.23461.0?token=9yStWppfKd3LiJBD2PII6BiNSHk%3D%401385723405%400
HTTP/1.1
> User-Agent: curl/7.15.5 (i686-redhat-linux-gnu) libcurl/7.15.5 OpenSSL/0.9.8b zlib/1.2.3 libidn/0.6.5
> Host: prod-se-03.ct.infn.it
> Accept: */*
>
< HTTP/1.1 403 Forbidden
< Date: Fri, 29 Nov 2013 11:02:42 GMT
< Server: Apache/2.2.15 (Scientific Linux)
< Vary: Accept-Encoding
< Content-Length: 347
< Content-Type: text/html; charset=iso-8859-1
```

Table 3.2 No access provided to URL with expired token or not coming from authorized IPs

3.2 Entering GridBox

The DPM WebDAV interface as it is, doesn't solve our problem, as a certificate or proxy is still required for use. We designed an architecture where a successfully authenticated user, with a Shibboleth session token, obtains a valid redirect link to be returned to his/her client so that there is direct access to the requested resource. This solution requires the introduction of a middle layer service that firstly authenticates the requesting user with the provided Shibboleth token, then secondly generates a proxy certificate and uses this proxy to forward the user client's request to the DPM WebDAV interface. This returns a redirection link, which gives the user client access to the requested file in the DPM disk node.

This mechanism worked if the client and the middle layer service reside in the same machine. But it stopped working when used by a remote client, resulting in an access forbidden error. The failure is caused by the fact that the redirection URL's token has been generated for the IP address of the middle layer server, therefore access is allowed only from that machine.

Our solution was to propose a modification in the protocol that brings the generation of the token on the DPM head node. We introduced an extra feature in the DPM authorization system, a "redirection delegation": a **delegator** client, with an authorized X.509 certificate which contacts the DPM head node WebDAV interface for a HTTP request, could add an extra parameter with the IP address of a **delegated** client. The head node, while generating the token, instead of using the IP address of the requesting **delegator** client, should use the IP address of the **delegated** client provided by the **delegator** client, if it has an authorized X.509 certificate. In this way, the redirection link, generated by the head node, could be forwarded to the delegated client that has the IP address presented by the delegator client, granting direct access to the disk node to complete the requested HTTP operation.

3.3 GridBox architecture

In the first implementation of this proposed solution, no WebDAV interface existed yet, and only a simple HTTPs/HTTP redirector was provided by the DPM developers on

top of DPM head and disk nodes¹⁰. We made slight changes to the source code of the DPM head node's redirector script by adding an extra parameter to the URL's query string, *authip*, with the IP address of the delegated client:

```

$ curl -k -i -E /tmp/x509up_u501 https://unict-diit-se-
01.ct.pi2s2.it/dpm/ct.pi2s2.it/home/cometa/TESTScienceGateway/riccardo/2011-11-
14/ciao.txt?authip=193.206.208.35
* Server certificate:
*   subject: /C=IT/O=INFN/OU=Host/L=DIIT UNICT/CN=unict-diit-se-01.ct.pi2s2.it
*   start date: 2011-05-25 13:11:58 GMT
*   expire date: 2012-05-24 13:11:58 GMT
*   subjectAltName: unict-diit-se-01.ct.pi2s2.it matched
*   issuer: /C=IT/O=INFN/CN=INFN CA
* SSL certificate verify result: self signed certificate in certificate chain (19), continuing anyway.
> GET /dpm/ct.pi2s2.it/home/cometa/TESTScienceGateway/riccardo/2011-11-
14/ciao.txt?authip=193.206.208.35 HTTP/1.1
> User-Agent: curl/7.15.5 (x86_64-redhat-linux-gnu) libcurl/7.15.5 OpenSSL/0.9.8b zlib/1.2.3
libidn/0.6.5
> Host: unict-diit-se-01.ct.pi2s2.it
> Accept: */*
>
< HTTP/1.1 302 Found
< Date: Fri, 29 Nov 2013 12:16:10 GMT
< Server: Apache/2.0.52 (Scientific Linux) mod_ssl/2.0.52 OpenSSL/0.9.7a mod_Gridsite/1.1.20
< Location: http://unict-diit-se-01.ct.pi2s2.it:777/gpfs/cometa/2011-11-
14/ciao.txt.692387.0?httpstoken=/dpm/ct.pi2s2.it/home/cometa/TESTScienceGateway/ric
cardo/2011-11-14/ciao.txt@212.189.144.115:GET:/gpfs/cometa/2011-11-
14/ciao.txt.692387.0:00000000:1385727380:/C=IT/O=INFN/OU=Personal%20Certificate/L=Ca
tania/CN=Antonio%20Calanducci/CN=proxy:unict-diit-se-01.ct.pi2s2.it:no-
token&httpsauthz=ej+IA0yKyXZwMGL091+3TYi6GA22ccDOptADIBuSS3W816e9BWgHPLAZdu
MVQPkQLJOyScj/qShCwIXfop3Fg/heACLpvCJG8bdStySmJ6DEGoWQvuc86CtGnEbi1ypZcUHNJ
/azOStN25Knfep7FbHeSkBh28GJqdlU7a3KMG4=
< Content-Length: 742
< Content-Type: text/html; charset=iso-8859-1

```

Table 3.3 Delegated download for first releases of DPM

¹⁰ <https://twiki.cern.ch/twiki/bin/view/LCG/DpmHttpsAccess>

The returned redirect URL will be only authorized within 3 seconds (this default is configurable) and can only be used only by the client that has the IP indicated in the *authip* query string parameter.

Since the release of EMI-2, DPM developers officially introduced a complete and standard WebDAV interface. Instead of patching the new release ourselves, we explained our requirements to the DPM team. The feature request was welcomed, including the support for the “redirection delegation” since LCGDM-DAV 0.13. Instead of using the query string to provide the head node with the IP address of the delegated client, we need to use the “X-Auth-IP” HTTP header in our delegator client requests. Additionally, for extra security, the X.509 Subjects (or DNs) of the allowed client delegators are listed in the Apache dav module’s configuration file, acting as an Access Control List.

So if we want to allow a client with IP address *193.206.208.201* to download a *demo.jpg* file, our delegator client, with a X.509 or proxy certificate whose DN is in the ACL of the head node, should issue the following command:

```
$ curl -k -v -E proxy --header "X-Auth-Ip: 193.206.208.201" -X GET https://prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg
> GET /dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg HTTP/1.1
> User-Agent: curl/7.33.0
> Host: prod-se-03.ct.infn.it
> Accept: */*
> X-Auth-Ip: 193.206.208.201
>
< HTTP/1.1 302 Found
< Date: Fri, 29 Nov 2013 17:44:40 GMT
* Server Apache/2.2.15 (Scientific Linux) is not blacklisted
< Server: Apache/2.2.15 (Scientific Linux)
< Link: <http://prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg?metalink>; rel=describedby; type="application/metalink+xml"
< Location: http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-29/demo.jpg.23461.0?token=Q4Ni2mOLduUCPvoD4Pr9UEDsxl8%3D%401385748080%400&dav_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg
< Vary: Accept-Encoding
< Content-Length: 485
< Content-Type: text/html; charset=iso-8859-1
```

Table 3.4 Delegated download with X-Auth-IP header

The Location URL returned, could be used now from the client with the given IP address to retrieve the file, without any X.509 certificate or proxy. This works seamlessly for upload operations using PUT requests, and all WebDAV operations using other HTTP verbs (PROPFIND, MKCOL, MOVE, COPY, etc).

Once we solved the authorization problem through the usage of the delegation header (*X-Auth-IP*), we were able to implement our delegator middle layer, named GridBox. We provided it with a RESTful interface whose task is to accept GET/PUT (and others) requests from a Shibboleth authenticated client. It then retrieves a valid proxy to forward the incoming requests to the DPM head node setting the *X-Auth-IP* header with the IP address of the Shibboleth authenticated and delegated client. The redirect Location URL, given by the DPM head node using a 302 Found status code, is returned back to the delegated client. So this, which is still listening on the connection to the GridBox endpoint, follows the returned Location URL, and is able to complete its GET/PUT (or others) operations directly (without a certificate) with the destination DPM disk node storage. Moreover, all the data flows in a direct data channel between the delegated client and the destination DPM storage, without any caching by middle server. The only data that transits on the GridBox middle layer are the few bytes of the request and response headers of the HTTP protocol. The complete GridBox architecture, with a sequence diagram, is illustrated in figure 3.2.

The following is an example of file download (GET) through the GridBox service. A valid Shibboleth session token is required:

```
$ curl -v -O -L -b
_shibsession_64656661756c7468747470733a2f2f676c6962726172792e63742e696e666e2e69
742f73686962626f6c657468=_cbe10d018a3051fbcd07dd2fce8872a1
https://Gridbox.ct.infn.it/vo.dch-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-
rp.eu/test/demo.jpg
* About to connect() to Gridbox.ct.infn.it port 443 (#0)
* Trying 193.206.208.35...
> GET /vo.dch-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg
HTTP/1.1
> User-Agent: curl/7.33.0
> Host: Gridbox.ct.infn.it
> Accept: */*
> Cookie:
```



```

_shibsession_64656661756c7468747470733a2f2f676c6962726172792e63742e696e666e2e69
742f73686962626f6c657468=_cbe10d018a3051fbcd07dd2fce8872a1
>
< HTTP/1.1 302 FOUND
< Date: Sat, 30 Nov 2013 11:16:59 GMT
* Server Apache/2.2.3 (Scientific Linux) is not blacklisted
< Server: Apache/2.2.3 (Scientific Linux)
< Content-Length: 613
< Location: http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-
29/demo.jpg.23461.0?token=1SFAvwfzu625uhwRDVQ3yV2IWOM%3D%401385811259%400
&dav_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg
< Connection: close
< Content-Type: text/html; charset=utf-8
* Closing connection 0
* SSLv3, TLS alert, Client hello (1):
} [data not shown]
* Issue another request to this URL: 'http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-
29/demo.jpg.23461.0?token=1SFAvwfzu625uhwRDVQ3yV2IWOM%3D%401385811259%400&d
av_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg'
* Connected to prod-se-03.ct.infn.it (193.206.208.163) port 80 (#1)
> GET /storage/vo.dch-rp.eu/2013-11-
29/demo.jpg.23461.0?token=1SFAvwfzu625uhwRDVQ3yV2IWOM%3D%401385811259%400
&dav_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg HTTP/1.1
> User-Agent: curl/7.33.0
> Host: prod-se-03.ct.infn.it
> Accept: */*
> Cookie:
_shibsession_64656661756c7468747470733a2f2f676c6962726172792e63742e696e666e2e69
742f73686962626f6c657468=_cbe10d018a3051fbcd07dd2fce8872a1
>
< HTTP/1.1 200 OK
< Date: Sat, 30 Nov 2013 11:17:39 GMT
* Server Apache/2.2.15 (Scientific Linux) is not blacklisted
< Server: Apache/2.2.15 (Scientific Linux)
< Content-Length: 8949
< Content-Disposition: filename="demo.jpg"
< Accept-Ranges: bytes
< Access-Control-Allow-Origin: *
< Access-Control-Allow-Methods: POST,GET,DELETE,PUT,OPTIONS,TRACE
< Access-Control-Allow-Headers: Content-Type,Content-Disposition,X-Requested-With,X-File-
Type,X-File-Name,X-File-Size
< Content-Type: image/jpeg
<
{ [data not shown]

```

Table 3.5 File download through GridBox

As shown in the previous curl command, no certificate or proxy has been provided, only a Shibboleth session token in a cookie.

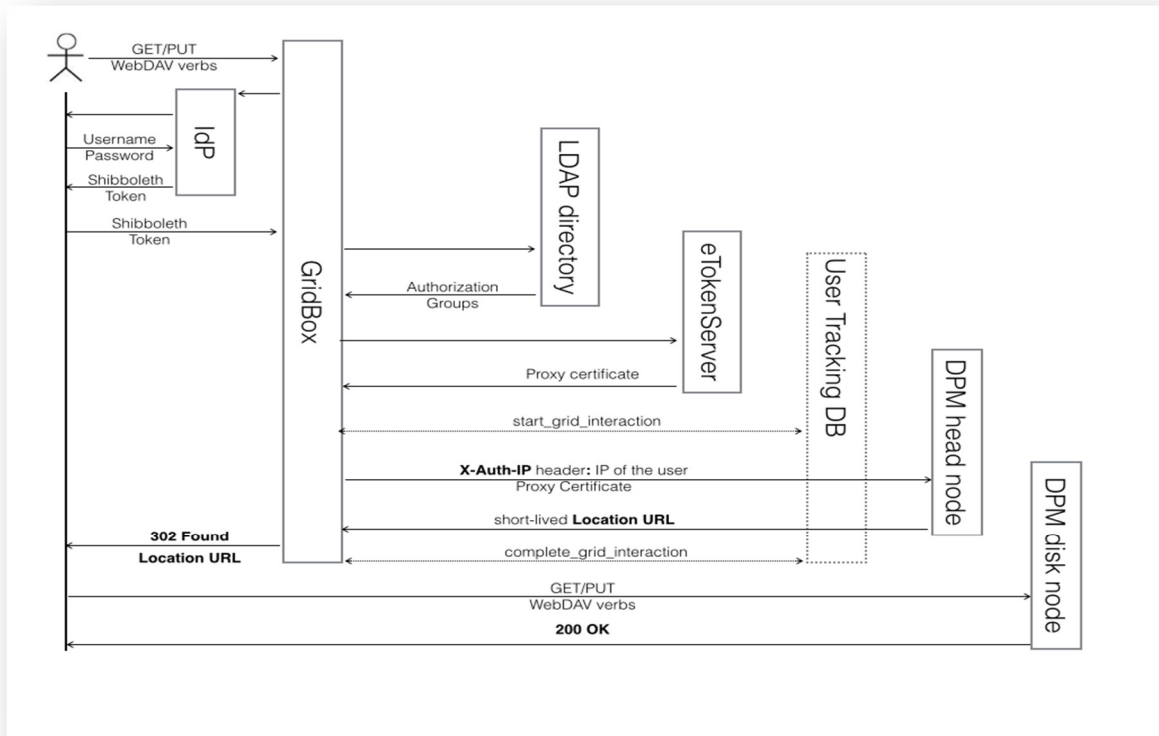


Figure 3. 2 GridBox Architecture, sequence diagram

3.4 Example of GET and PUT operations

Following the diagram illustrated in Figure 3.2, consider all the interactions involved to request a file download from a web browser:

- A client issues a GET request to the GridBox server endpoint. The Virtual Organization is provided, the destination Grid storage element, and the full path of the file it would like to download:
 - <https://<Gridbox host>/<vo>/<storage>/<path>>
 - For example: <https://Gridbox.ct.infn.it/vo.dcb-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dcb-rp.eu/test/demo.jpg>

- As the REST APIs are Shibboleth protected, our GridBox server, acting as a Shibboleth Service Provider (SP), starts the authentication process. It redirects the user's client to the configured Discovery Service (DS) (not shown for simplicity) where the user chooses the Identity Federation they belong to. This brings the user to another page from where they choose their institutional Identity Provider (IdP). Finally the user authenticates with a username and password, or any other authentication mechanism provided by his/her IdP. If the authentication phase is successful, a Shibboleth session token is generated and returned to our client (web browser) that stores it in a cookie.
- A subsequent redirect will bring our client back, with the Shibboleth token to the GridBox service. This is accepted and the requested operation is completed;
- The GridBox service can now inspect the HTTP server environment variables (it runs as a WSGI application in an Apache child daemon), to extract the SAML attributes provided by the IdP that has authenticated the user. For our implementation we currently use the *mail* of the user as a key to verify the authorization in our LDAP directory service. This allows GridBox to retrieve the user's groups and roles (not used at the moment). Moreover it also retrieves the client's IP address (from REMOTE_ADDRESS) to be delegated.
- If the user is authorized, the GridBox service needs a proxy certificate to establish a connection with the DPM head node and forward the request coming from the client. This proxy is obtained by making a REST call request to the API of an eTokenServer, together with the name fo the Virtual Organization the user belongs to. The eTokenServer generates and returns a proxy certificate authorized for the given VO from a set of Robot Certificates stored on hardware tokens. The eTokenServer uses an ACL based on IP addresses, and can be contacted only by a list of authorized machines. Currently, a new proxy is requested every time a new request comes in. Caching will be implemented to reuse not-expired proxies and avoid to avoid to making unnecessary HTTP requests.
- Once a proxy is retrieved and stored in the GridBox server's filesystem, we need to log the upcoming transfer or access operation. This is a mandatory requirement for European Grid Infrastructures when Grid operations have not been directly authenticated via a X.509 certificate. We record in a relational DB several fields related

to the user: their name, surname, email address from the IdP, kind of operation, the source IP, the destination storage element and the path of the file they are accessing. This is stored in the *active_Grid_operations* table;

- The GridBox service is now ready to forward the client requested operation (GET) to the DPM head node. An extra HTTP header is added to the request, *X-Auth-IP*, with the IP address of the user's client, and authenticating the connection with the proxy returned by the eTokenServer.
- Supposing that the request path exists in the name server of the DPM head node (DPNS), and the proxy certificate maps to a DPM user authorized to read the given file, the DPM head node builds the short lived token, authorized only for the *X-Auth-IP* address. It appends the token as a parameter in the query string of the URL to access the requested file in the DPM disk node file system, and returns back to the GridBox service the redirect link in the response Location header containing a "302 Found" status:
 - **Location:** http://prod-se-03.ct.infn.it/storage/vo.dch-rp.eu/2013-11-29/demo.jpg.23461.0?token=1SFAvwfzu625uhwRDVQ3yV2lWOM%3D%401385811259%400&dav_sfn=%2Fdpm%2Fct.infn.it%2Fhome%2Fvo.dch-rp.eu%2Ftest%2Fdemo.jpg
- The GridBox service, still connected to the client, then forwards back the very same response back to the requesting client (with another 302 Found status).
- The GridBox service contacts again the User Tracking services, moving the record previously added from the *active_Grid_operations* table to the *completed_Grid_operations* table;
- The user's client finally follows the redirect found in the GridBox service response and sends a direct GET request to that URL.
- The DPM disk node, will verify if the token has expired and if the IP address of the requesting client is authorized. It then decodes the token parameter, and if successful, finally sends the file data directly to the client returning a 200 OK status code.

3.5 GridBox file transfer REST APIs

The following APIs are currently available to interact with the GridBox data management service for file transfers:

- URL: https://<Gridbox_host>/<vo>/<se>/<path>
Method: GET
- URL: https://<Gridbox_host>/<vo>/<se>/<path>
Method: POST
- URL: https://<Gridbox_host>/<vo>/<se>/<path>
Method: PUT

The last two APIs return a JSON object, with the following structure, in case of success:

```
{
  "status": <resp_status>,
  "redirect": <redirect_url>
}
or:
{
  "status": <resp_status>,
  "reason": <resp_reason>,
  "response": <response_text>
}
```

in case of a failure.

Table 3.6 JSON response of GridBox REST APIs

In the case of an upload, either via a POST or PUT method, the GridBox service does not currently return a 302 Found redirect, so that the client could automatically follow the redirect. Instead, it returns a 200 OK status code and the redirect URL will be contained in the JSON object described above. It is now responsibility of the client to immediately issue a subsequent POST or PUT request, according to its capabilities (some older web browser can only handle POST operation for uploads, using a multipart/form-data enctype), again, before the time validity of the token expires. We chose this approach because the automatic redirection of PUT or POST requests is actually forbidden by the RFC¹¹ defining the standard of the HTTP 1/1 Protocol:

```
''The action required MAY be carried out by the user agent
without interaction with the user if and only if the method
used in the second request is GET or HEAD''
```

Table 3.7 RFC 2616 on PUT redirects

The automatic redirection is allowed only for GET or HEAD requests.

Moreover, not so many clients are able to follow a PUT or POST redirect. In any event the RFC does not forbid manual redirection. Some clients actually open a dialog box informing the user and requesting permission to proceed. Cyberduck¹² a very popular open source and FTP, WebDAV and S3 client, has accepted a feature request from the DPM developers and ourselves to implement the PUT redirection with user intervention, as described in this ticket¹³.

¹¹ RFC 2616, <http://tools.ietf.org/html/rfc2616#page-61>

¹² <http://cyberduck.io/>

¹³ <https://trac.cyberduck.io/ticket/6586>

3.6 Obtaining Shibboleth sessions tokens for non browser clients

So far we have described users accessing GridBox services via a Web browser through the REST URLs or via a Web application or portal. Obtaining a token is transparent to them, at the end of the authentication process. SAML and all of its implementations, Shibboleth included, have been designed to provide authentication to the Web and to be used from a browser, due to its redirection protocol. However, SAML includes an Enhanced Client or Proxy profile¹⁴ for authentication that is designed for clients other than browsers, such as desktop applications, server-side code running in web environments. It is designed to work with HTTP and assumes a session-based security approach between the client and the HTTP server, usually via a cookie. Recently a new version of the ECP profile¹⁵ has been standardized.

As one of the requirements for accessing GridBox services was from mobile devices, but not from their built-in browsers, where the classical mechanism would be used to get a Shibboleth token. Rather we wanted to be able to access GridBox services from native applications and we needed a way to obtain a valid session token. Our solution was to use a hybrid app. Hybrid apps are built using HTML5 and native device APIs. HTML5 and CSS are running inside a WebView, a component embedded in a native application, capable of rendering HTML and CSS, and eventually running JavaScript code. Platforms that use this approach, for example PhoneGap or Adobe Cordova, provide a way to establish a bridge between native and the WebView context. Having a WebView component in our app allows the user to be directed to the Discovery Service first and to their Identity Provider later, completing the authentication process in this embedded browser. Once the IdP page has been submitted, we intercept the response returned by the Service Provider (GridBox in our case) and read its cookies to look for a valid Shibboleth session. If found, we save this and use it to send all the following requests to our Shibboleth protected service. Once the token has expired, we use the WebView again and request the user to re-authenticate. This mechanism is very similar to the approach used by many mobile native apps that use the OAuth2 protocol for authentication and authorization purposes.

¹⁴ ECP, <https://wiki.shibboleth.net/confluence/display/SHIB2/ECP>

¹⁵ 2.0, <http://wiki.oasis-open.org/security/SAML2EnhancedClientProfile>

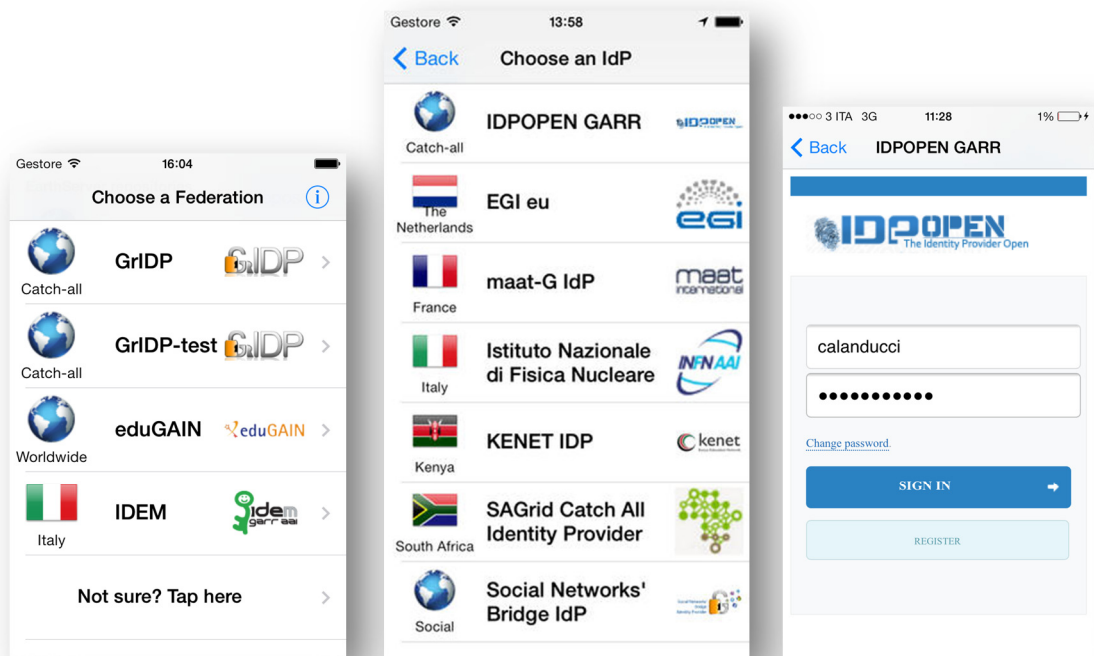


Figure 3. 3 GridBox native iOS client – Shibboleth authentication

3.7 WebDAV GridBox interface

GridBox file transfer RESTful APIs only allow the basic operations such as uploads and downloads of files from the command line, from ad-hoc Web Applications and mobile applications. We therefore thought that dealing with very large sets of (and big) files or managing the namespace of Grid storage services could be better handled using desktop clients, or even better the built-in tools of desktop operating systems. Our reference Grid storage implementation, DPM, recently introduced a full WebDAV interface, as previously mentioned. WebDAV clients are now included in the recent releases of operation systems, Windows since Windows XP SP3 with Web Folders, Mac OS and Linux distribution with Gnome or KDE desktop environments. Being able to “mount” a remote Grid storage directory on the user’s desktop file system is another goal of our study in the effort of improving the accessibility of Grid based storages.

We extended our GridBox service to provide a WebDAV interface, adding one extra API which is capable of handling some of the extended HTTP operations of WebDAV, such as PROPFIND, MKCOL, MOVE, COPY. This API has the following endpoint:

- https://<Gridbox_host>/shibdav/<vo>/<se>/<path>

According to the method used, extra parameters are required into specific request headers¹⁶.

This API forwards client's requests to the DPM head node, that manages the filesystem namespace of a given storage server. It authenticates with a proxy retrieved from the eTokenServer, logging the operations, and returning back to the client the XML response containing the properly modified paths from the DPM head node. Those modifications are required to map the path structure provided in the API to the shorter DPM namespace that does not have /shibdav/<vo>/<se>/ part.

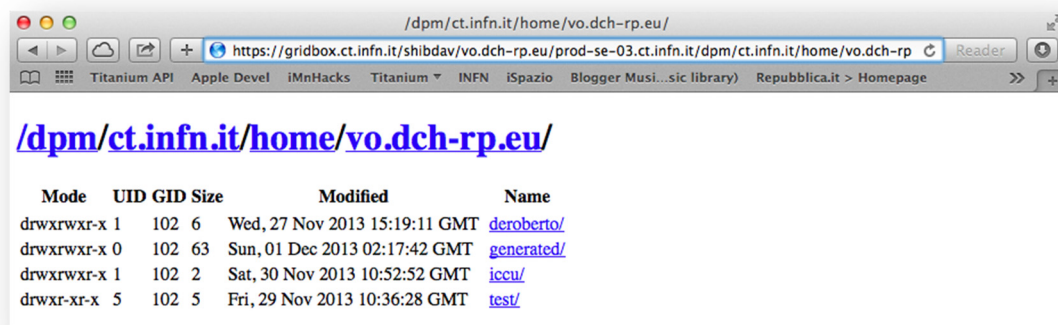


Figure 3. 4 Accessing the Shibboleth protected WebDAV interface from a browser

However, the main concern here is related to the authentication mechanism. WebDAV does not explicitly require a specific type of authentication, but in practice all the available clients generally offer Basic and Digest Authentication only. Those authentication mechanism are quite weak cause they pass the credentials over HTTP headers.

¹⁶ <http://www.webdav.org/specs/rfc4918.html>

Some work has been already done to integrate WebDAV with Shibboleth authentication mechanisms. This is technically possible and one can just install the Service Provider software into the WebDAV server and protecting the server endpoint with a valid shibboleth session (as it has been done in the previous REST API). In practice this requires a way to pass in a Shibboleth token to the server, and providing some sort of embedded Web browser to authenticate. None of the clients currently available both built-in in OSes and stand alone, provides such a feature. One solution investigated by [RRX11], was to develop a Java Standalone client, on top of the open source *Sardine* project, that shows a dialog window to let users input their credentials to their IdP, manipulating and processing under the hood SAML responses and submitting the SAML POST profile.

This approach however would not allow our users to use their OS built-in clients or the popular desktop WebDAV clients. Other work to solve this problem has been done by EDINA and the Data Library division of Information Services at the University of Edinburgh, in the context of the WSTIERIA project¹⁷ and described here [WST10a, WST10b]. They proposed separating the authentication flow from the client. It starts running in a web environment, and after successful authentication, the web page returns a URL with a valid token, identifying the authenticated user, to be used in any WebDAV client or desktop application.

Our solution is a slight modification of this, using the Shibboleth token as part of the URI:

<http://<Gridbox host>/<shibboleth session token>/>

This Shibboleth token is obtained by first visiting a Shibboleth protected endpoint of GridBox:

<https://<Gridbox host>/webdav/<vo>/<se>/<root path>>

If the authentication is successful, it returns the URL shown above with the Shibboleth session token to be used in any WebDAV client. The `<root_path>` parameter represents the starting folder from where the client has access in the storage name space. If the returned URL, is used in one of the OS built-in WebDAV client, this will allow the user

¹⁷ Web Service Tiered Internet Authorization, http://edina.ac.uk/projects/wstieria_summary.html

to browse and copy files to the given Grid storage as if it is a local disk, as showed in the picture below.

As a security mechanism, the returned WebDAV endpoint can be used only by the same machine from which the authentication request has been issued. This is internally implemented by checking if the WebDAV client's IP is the same of the Web Browser client used in the authentication stage. The WebDAV URL will have the very same expiration time of the Shibboleth token.

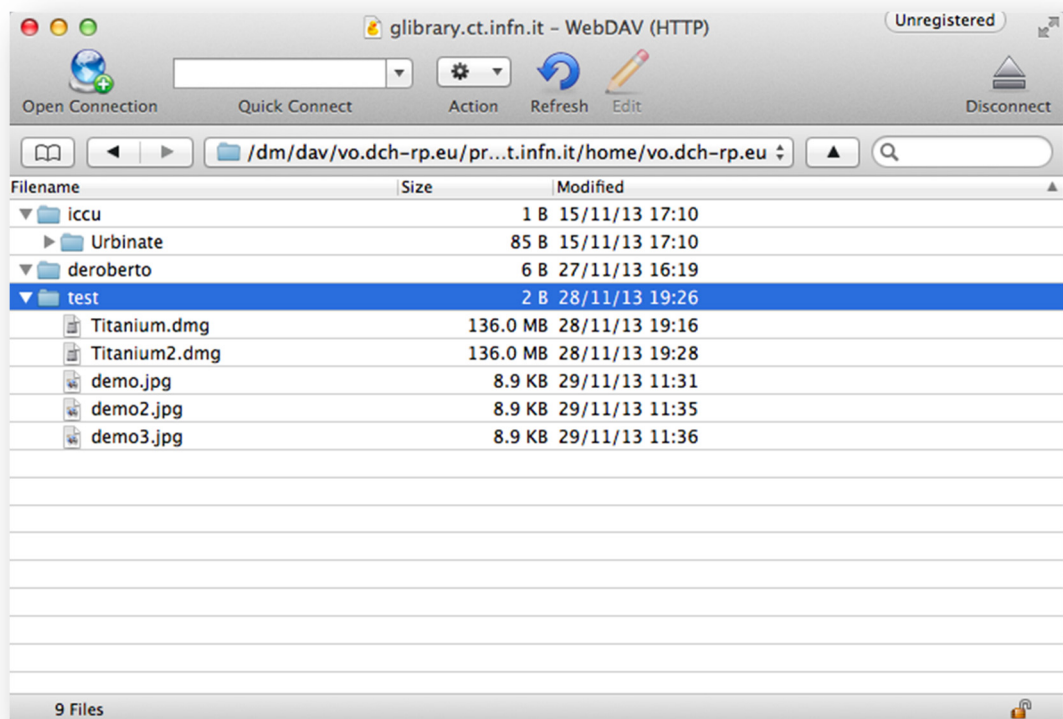


Figure 3. 5 browsing the a Grid storage using a desktop WebDAV client

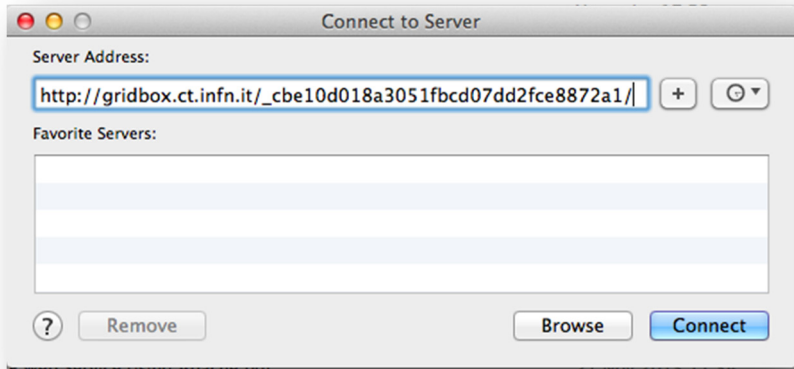


Figure 3. 6 using the built-in WebDAV client built-in in Mac OS to mount a Grid storage directory

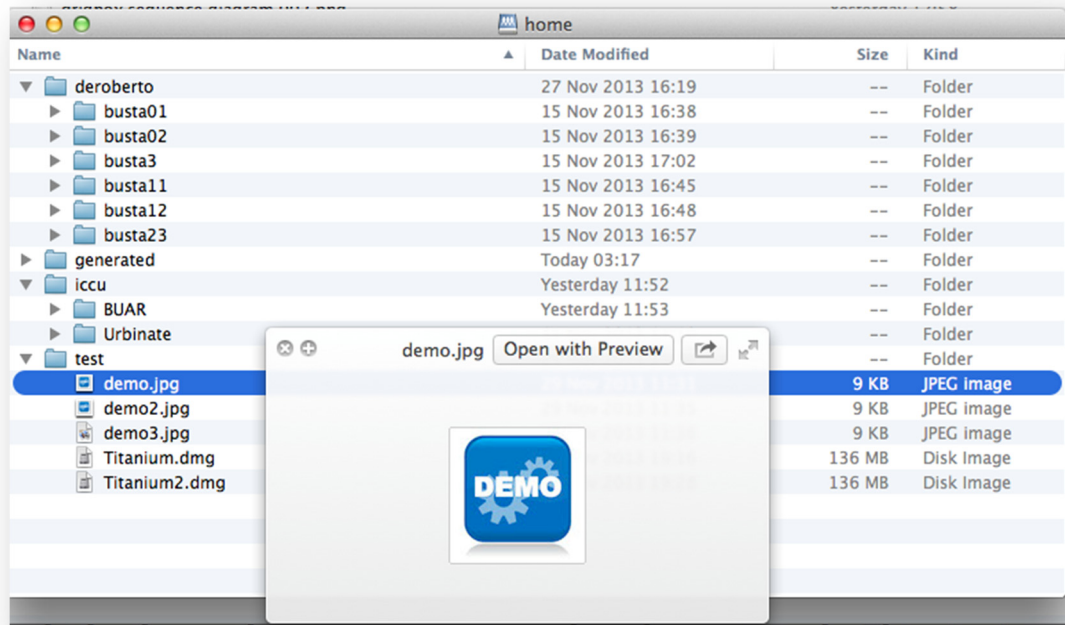


Figure 3. 7 browsing the Grid storage through Gridbox as a local disk

3.8 GridBox file sharing

One extra requirement is often requested from user communities: the ability to share files amongst them, using public links, a feature provided by many Dropbox-like services. While we could easily implement this functionality on top of GridBox architecture the

scientific communities, that are the target users of our infrastructures, require added security. Following the delegation approach at the core of GridBox, the service can create public URLs that could be used only for a given client (IP), a given number of times and for a limited amount of time.

Our Shibboleth protected “public” URL creation REST API, looks like this:

- https://<Gridbox_host>/public/<ip>/<counts>/<exp_time>/<vo>/<se>/<path>

For example:

- <https://Gridbox.ct.infn.it/public/193.206.208.201/3/1387879382951/vo.dch-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg>

The previous file could be downloaded only by clients coming from 193.206.208.201, for a maximum of three times, by Dec, 24th of 2013, 11:03:02 (GMT + 1). If for some reason, complete public access should be allowed, it's possible to use 0.0.0.0 as IP address and 0s both the count and expiration time, after a proper configuration of the service.

This API returns a plain HTTP URL with the following format:

- http://<Gridbox_host>/dl/<guid>/<filename>

3.9 Execution time comparison

Finally, we have conducted tests to verify how much overhead is generated by the delegation and redirection mechanisms introduced by the GridBox service. We have made a simple download test, retrieving a small file of 9 KB (demo.jpg, 8.949 bytes) and a larger file of 3,2 MB (Avanguardia1911_173_0002.jpg, 3.152.252 bytes).

As expected, the overhead was really minimal even for the small file, in the order of milliseconds.

```

$ time curl -O -L -k -E x509up_u507 https://prod-se-
03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg
% Total % Received % Xferd Average Speed Time Time Time Current
           Dload Upload Total Spent Left Speed
100 485 100 485 0 0 999 0 --:--:-- --:--:-- --:--:-- 1002
100 8949 100 8949 0 0 13358 0 --:--:-- --:--:-- --:--:-- 13358

real    0m0.680s
user    0m0.010s
sys     0m0.007s

```

```

$ time curl -O -L -k -E x509up_u507 https://prod-se-
03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-
rp.eu/iccu/BUAR/RML0067972/1911/IMG/USAGE2/Avanguardia1911_173_0002.jpg
% Total % Received % Xferd Average Speed Time Time Time Current
           Dload Upload Total Spent Left Speed
100 567 100 567 0 0 1106 0 --:--:-- --:--:-- --:--:-- 1109
100 3078k 100 3078k 0 0 791k 0 0:00:03 0:00:03 --:--:-- 1003k

real    0m3.900s
user    0m0.035s
sys     0m0.077s

```

Table 3.8 File download requests directly from a Grid storage using a local X.509 certificate

```

$ time curl -O -L -b
_shibsession_64656661756c7468747470733a2f2f676c6962726172792e63742e696e666e2e69
742f73686962626f6c657468=_35ca961086b532d76d08ec23e810e33a

```

```
https://Gridbox.ct.infn.it/vo.dch-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/test/demo.jpg
```

```
% Total   % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total  Spent    Left     Speed
0  613    0   0   0   0   0   0  0  0  0  0  0  0  0  0  0  0  0  0
100 8949 100 8949  0   0 13306  0  0  0  0  0  0  0  0  0  0  0  0 13306
```

```
real 0m0.683s
user 0m0.014s
sys 0m0.008s
```

```
$ time curl -O -L -b
```

```
_shibsession_64656661756c7468747470733a2f2f676c6962726172792e63742e696e666e2e69742f73686962626f6c657468=_35ca961086b532d76d08ec23e810e33a
```

```
https://Gridbox.ct.infn.it/vo.dch-rp.eu/prod-se-03.ct.infn.it/dpm/ct.infn.it/home/vo.dch-rp.eu/iccu/BUAR/RML0067972/1911/IMG/USAGE2/Avanguardia1911_173_0002.jpg
```

```
% Total   % Received % Xferd  Average Speed   Time    Time     Time  Current
           Dload  Upload  Total  Spent    Left     Speed
0  781    0   0   0   0   0   0  0  0  0  0  0  0  0  0  0  0  0  0
100 3078k 100 3078k  0   0  779k  0 0:00:03 0:00:03 0:00:00 929k
```

```
real 0m3.959s
user 0m0.038s
sys 0m0.078s
```

Table 3.9 File download request through GridBox with a Shibboleth session token

	9KB File	3.2MB File
Direct with proxy	0m0.680s	0m3.900s
GridBox with Shib token	0m0.683s	0m3.959s

Table 3.10 Execution time comparison

Chapter 4

GLIBRARY, A DIGITAL ASSET MANAGEMENT SYSTEM FOR GRID

In every Grid infrastructure a huge amount of distributed storage resources is available to save user's data. However, few tools are provided by Grid middleware to easily search and retrieve files a user is looking for, because of the weakness of additional information describing file contents. File Catalogs offer virtual file systems distributed among Grid Storage servers, that can help to organize data in hierarchical structures, grouping them according to some user's criteria, but they do not provide a way to describe file contents. On the other hand, Grid Metadata services can be used to attach additional information to files, but they are usually quite complex to use, and then not suitable for novice users.

On top of the GridBox service presented in the previous chapter, we designed the gLibrary platform, aiming at offering a secure and easy-to-use system to manage digital assets stored on a distributed Grid infrastructure. This goal can be achieved orchestrating a set of Grid services hiding their interactions with a proper business logic, and providing intuitive front-ends, accessible from everywhere. Users do not have to care about or know the complexity of the underlying infrastructure and the geographical dislocation of their data and they can consider the available Grid storage as a huge virtual disk.

4.1 Features

gLibrary can be used to store, organize, search and retrieve any kind of digital assets in a Grid environment. By digital asset, we refer to any kind of digital file (also referred as digital object) and its associated metadata. It can be employed by several categories of users

that need a secure way to save, manage and share their assets. Exploiting the fine grained authorization system offered by Grid infrastructures, data providers can define permissions on their assets in order to grant or deny access to given users, groups or even whole organizations.

A gLibrary user with the role of administrator can create and manage one or more repositories. Every repository organizes its assets by *Types* and/or *Collections*. A *Type* is a list of metadata attributes that describe a specific kind of asset to be organized by the system.

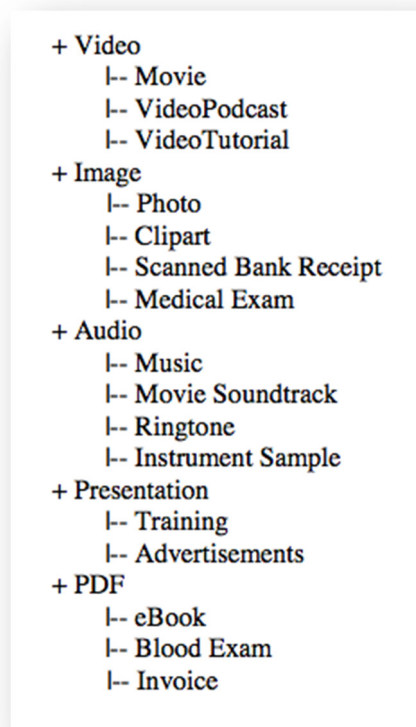


Figure 4. 1 Example of types' tree

Other systems, based on the MVC architecture, call it *Model*. All the assets of a given type can be later queried by one or more of those attributes. Each type can then have multiple subtypes with additional attributes set. All types share a common attribute lists (defined by the repository administrator, generally the Dublin Core Metadata Schema set is chosen). So for every repository a hierarchy of types is defined, before users can start uploading assets. Figure 4.1 shows an example of types' hierarchy and figure 4.2 present an example of types and their attributes that can be used for a repository of multimedia assets.

Type	Attributes' list
Audio	Format, Bitrate, Samplerate, Time
Music	(Format, Bitrate, Samplerate, Time), Name, Artist, Album, Genre, Tracknumber, Year, Artwork, Lyric, Rating
Presentation	Format, NumOfPages
Training	(Format, NumOfPages), Title, Runtime, Speaker, Author, Subject, Event, Date, Type
(Root)	FileName, SubmissionDate, Description, Keywords, LastModificationDate, Size

Figure 4. 2 Example of Types and Attributes' list

Assets are associated with the proper type in the registration/upload process. An asset cataloged as a given subtype inherits the attributes of its parent type. Of course, types will be defined according to the users' needs and taking into account the assets they want to manage. The flexibility and extensibility offered by this type system allow different communities to adopt gLibrary for any cataloguing purpose.

When uploading a file, users have to link this asset with a suitable type, filling its attributes. Input files can be read from local disks, network shares, HTTP/FTP servers, and replicated to one or more Storage Servers on which the user is authorized. gLibrary can also manage assets already present on Grid resources, through direct access to File Catalogs to discover the physical location of files. No data copy is performed in this case, instead users only have to select the right type and set the required attributes.

gLibrary servers can host multiple repositories, which can have their own hierarchy of types and can be accessed by different users. Moreover, types and related entries are only visible by users with appropriate privileges. □ When the user is properly authenticated and has chosen one of the available repository on which he/she has access, the asset browser interface will show a library types' tree (Fig. 4.3). By choosing one of the available types, the list of all the assets of the selected type will be shown, together with the value of their attributes. This list can be sorted by any of the available attributes and a cascading set of per-type defined filters can be applied to easily find the assets the user is looking for.

When selecting one attribute of the current type as a filter, a separate list of all the possible distinct values of the chosen attribute will be displayed.

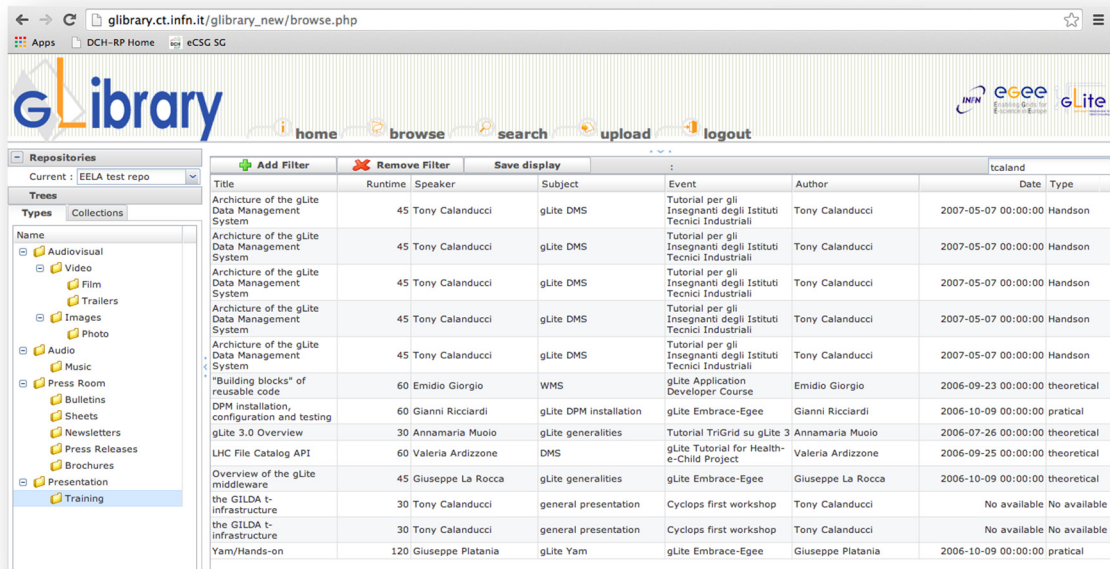


Figure 4. 3 Test repository for the EELA project – hierarchy of types on the left

One or more values of this list can be selected to filter the asset result set. A second filter can be applied, again selecting one different attribute, that will generate another list of all the distinct attribute's values. While generating the second list, the choices made in the first filter's list will be taken into account, showing only the values that pass all the filters. Further filters can be applied with the same mechanism, thus allowing a dynamic reduction of results. A similar filtering system can be seen in the iTunes software to organize iPod/iPhone/iPad music libraries, where the lists of all the Genres, Albums, Singers of the iPod owner songs are shown. This mechanism was improved in gLibrary, so that users can choose to apply any number of attributes as filters and their order of application. This intuitive and fast browsing system represents one of the unique feature of gLibrary and allows to find assets very quickly even if their number is huge. Figure 4.4 shown a sample of attribute filters used to browse training presentations.

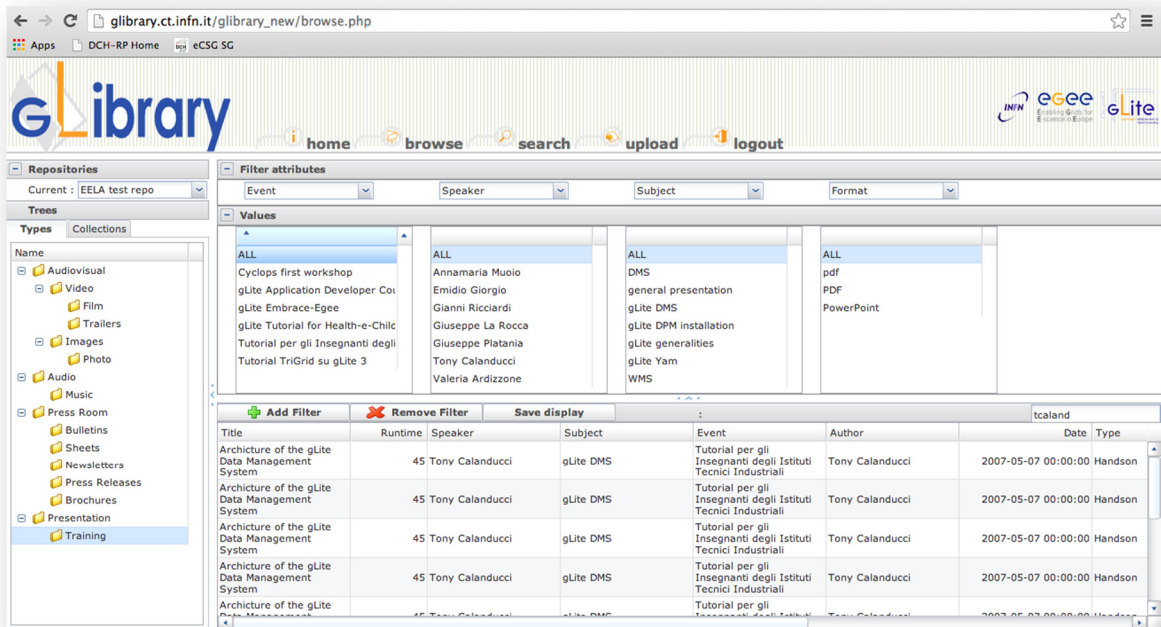


Figure 4. 4 Test repository – Available filters for the Training type

Assets can also be grouped by *collections*, allowing to:

- organize entries belonging to different types. For instance *Projects*, where each *Project Name* is a category that collects all project related assets; or working *Teams*, where each *Team* is a category.
- define subsets of assets belonging to the same type. Some examples could be *Playlists*, where each *Playlist Name* represents a category, grouping some assets of the type Music or the creation of a *My Favorites Movies/Songs* category.

Collections have no specific attributes but they are also organized in a hierarchical way. They are useful to define new views for users' assets. Unlike types, which are defined by the administrator only, user collections can be created by the users themselves. Assets can belong to only one type but they can be tagged with many collections.

We can also set permissions on collections using a fine-grained authorization mechanism. Each asset, type and collection has a set of ACLs (Access Control Lists) that restricts its usage, allowing asset owners to grant access to a whole organization, selected

groups of users or just a single user. Those entries, types and categories on which users do not have permissions, are also not visible from the browsing interface.

Three users' roles have been defined in gLibrary:

- *Repository Managers* – library administrators, which define type hierarchies, attributes, filters and permissions. They also grant or deny upload rights to the users;
- *Data Providers* – users allowed to upload/register assets and edit their attributes on libraries for which they have rights. They can also set ACLs on their assets and manage categories' hierarchy.
- *Generic Authorized Users* – users that can browse and download assets for which they are authorized. They have read-only access.

Technically we could have granted access to guest users, but EGI policies require that users accessing data on Grid resources need to be identified and tracked.

4.2 Interaction with gLite/EMI Grid services

The first implementation of gLibrary, started 3 years ago, was based on the service provided by the EGEE gLite Grid middleware. Born from the collaborative efforts of more than 80 people in 12 different academic and industrial research centres as part of the EGEE Project, gLite provided a bleeding-edge, best-of-breed framework for building Grid applications tapping into the power of distributed computing and storage resources across the Internet. Most of the services developed in the context of the EGEE project have been evolved and integrated into the EMI middleware, so we were able to move gLibrary from gLite to EMI without big efforts. As a set of lightweight services to implement a geographical distributed Computational Data Grid, gLite/EMI seems perfect to implement and deploy the system we are presenting. In particular, gLibrary, being a data-oriented platform, makes use of the following gLite/EMI services:

1. Storage Elements (SEs) provide uniform access to data storage resources. They can be single disks, large disk arrays or tape-based Mass Storage Systems. Our deployment use DPM storages, as described in the previous chapter

2. AMGA Metadata Catalog stores metadata describing the contents of Grid files, allowing users to search for entries based on their descriptions.
3. LCG File Catalog (LFC) maps logical filenames to the physical locations of replicas of a file stored in one or more Storage Elements.
4. Virtual Organization Membership Service (VOMS) [ACC+04]: a service that allows a detailed definition of users' privileges and roles according to abstract entities called Virtual Organizations (VOs), which group users, institutions and resources in the same administrative domain.
5. Information Service (IS) provides information about Grid resources and their status. In particular, IS is used to discover the available SEs for a given VO.

4.3 Architecture

gLibrary assets are saved as one or more replicas on Storage Elements. The actual SE implementation in use on the EGI infrastructure, based on disks is Disk Pool Manager (DPM), allows to set fine-grained permissions on files with owners, groups and ACLs, satisfying gLibrary security constraints.

The AMGA Metadata Catalog has a central role in the gLibrary architecture. It is the repository of all the assets' attributes, types and collections hierarchies, available libraries and physical file locations. Thanks to its powerful authorization system, it allows permission management on assets, types and collections, ensuring that the required information are accessible only to authorized users. User-defined groups and users' roles (Repository Managers, Data Providers and Generic Users) are also handled by AMGA.

Users' requests, such as browsing/searching of/in the libraries, are converted into SQL-like queries and sent to the AMGA backend. As some benchmark studies on performance demonstrated [SK06b], AMGA's overhead is very low and this guarantees short response times by the gLibrary front-ends.

If a user wants to add metadata information to an existing file, he/she can use gLibrary interfaces to browse the File Catalog of the VO he belongs to, and find the location of its replicas in Grid SEs. After selecting the file and choosing its type from the existing

hierarchy, the user can add further information by filling out a form generated dynamically according to the type selected.

The Information System is queried by the user to find out which Storage Elements can be used to upload his/her assets.

Authentication and authorization in gLibrary has evolved along the years. In the first releases of gLibrary, Virtual Organization Membership Services (VOMS) has been used for authentication and authorization purposes and user role mappings. This meant each gLibrary user needed to have to its own a X.509 Grid certificate and to belong to a Virtual Organization. In order to be authenticated directly as an AMGA user, a proxy certificate – containing information about VO membership and role – was required. Grid resources usage policies are regulated on VO basis. This means that a user could only upload his/her assets to SEs which accepts members of that VO. Moreover, VOMS roles were needed to be able to administer libraries as gLibrary Repository Manager. Thus, the assignment of this role was demanded to the VOMS managers.

Recently gLibrary has been integrated with the GridBox service, so this means that X509 certificates or proxies are no more required. Users are then authenticated via their federated identity, using a Shibboleth session token, as described in chapter 3, while authorizations are currently managed both in an external LDAP directory and into AMGA service, with its powerful group-based and ACL-based system. The LDAP directory is used to grant or deny users access to specific repositories according to their membership to Virtual Organizations or projects, while the internal AMGA user management defines permissions and roles within each repository, for types, collections and entries.

As a dependency from GridBox, gLibrary makes use of the eToken Service and the User Tracking DB. It uses the eToken Service to generate proxies from the robot certificate of the project/VO the users belongs to. The User Tracking DB is again used to log user details, coming from the SAML attributes of the Shibboleth token, and the requested Grid operations.

4.4 gLibrary first Implementation

gLibrary implementations have changed along the years, following the changes of the architecture and the evolution of the authentication mechanism.

The first prototype of gLibrary was implemented as a web application. The front-end was developed as a Web 2.0 application to offer a pleasant user experience. This implies the usage of technologies like AJAX, Javascript and Java Applets to offer a desktop-like interface. On the server machine, we deployed Apache 2.0 with the PHP 5 module. The business logic of the application is built as PHP 5 application that interacts with several Grid services. Besides the PHP-enabled web server, the machine runs the AMGA Metadata Catalog with its PostgreSQL backend and the command line tools to interact with the VOMS servers, Information System, LFC File Catalogs and DPM Storage Elements. Since AMGA is the most used Grid service by this web application, it has been decided to develop a low level PHP 5 API whose role is to interact directly with the AMGA server daemon. Figure 17 summarizes the web front-end architecture.

Authentication on Grid services made use of VOMS-enabled proxy certificates. These were created starting from X.509 personal public-private user key pairs. As it's considered a deprecable practice to send them over an Internet connection, gLibrary created a local proxy in the user's machine. This has been achieved writing a Java applet that uses Globus CogKit APIs to generate a local proxy without VOMS extensions. These were added once the proxy is sent over HTTPS to the web server, where VOMS command line tools were called to add VO Membership and Roles to the user's proxy. Then, the web application could be authenticated on all the necessary Grid services and interact with them on behalf of the user. Proxies were destroyed when the user signs out from the system.

Interactions with the SEs were done using their Storage Resource Management (SRM) interfaces, guaranteeing standard compliance with Grid Storage management and decoupling gLibrary from any specific Storage System architecture.

To avoid double transfers of the uploaded/downloaded files from the user's machine to the web server and then from the web server to the chosen Grid Storage Element and vice versa, another Java applet was developed to move files directly between user's client and Grid Storage Elements. Transfers were carried over the GridFTP protocol, a high-performance, secure, reliable data transfer protocol optimized for high-bandwidth wide-area networks. Once again, GridFTP APIs were provided by the Globus CogKit.

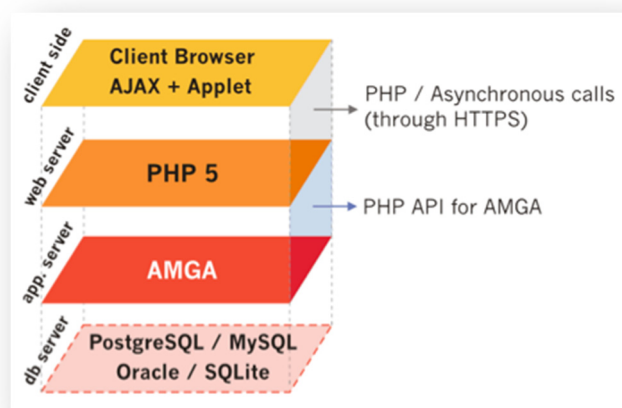


Figure 4. 5 first prototype of gLibrary implementation

4.5 gLibrary implemetation with GridBox and Federated authentication

In the last couple of years, gLibrary moved from a simple web application to a client-server architecture. We decoupled the backend from the front-end and we make gLibrary metadata and file transfer services accessible through a set of REST APIs. One of the biggest changes was the introduction of the federated authentication, based on SAML. So, besides Grid authentication based on X.509 certificates and proxies, users can now access gLibrary managed repositories being authenticated with their institutional credentials.

The REST APIs expose the access to the abstractions provided by gLibrary like repositories, types, collections, entries, filters, hiding all the Grid-related intricacies behind

the scene. Data movements APIs are provided by GridBox, so the Java Applets, used by the previous web front-end, were unnecessary anymore to achieve direct transfers between the user's client and storage servers.

gLibrary core services and REST APIs are now implemented using Python. REST APIs are deployed as a WSGI module¹⁸ in an Apache container. gLibrary metadata service has been developed using the Django framework, while all the file transfer GridBox APIs have written using a lightweight and smaller Python microframework, Flask. All the metadata APIs returns JSON data, while GridBox WebDAV ones returns XML according to the protocol.

Finally, an OAI-PMH interface has been implemented on top of the gLibrary metadata services, to allow external harvesters the extraction of gLibrary repositories' metadata. This endpoint has been developed with the MOAI libraries¹⁹.

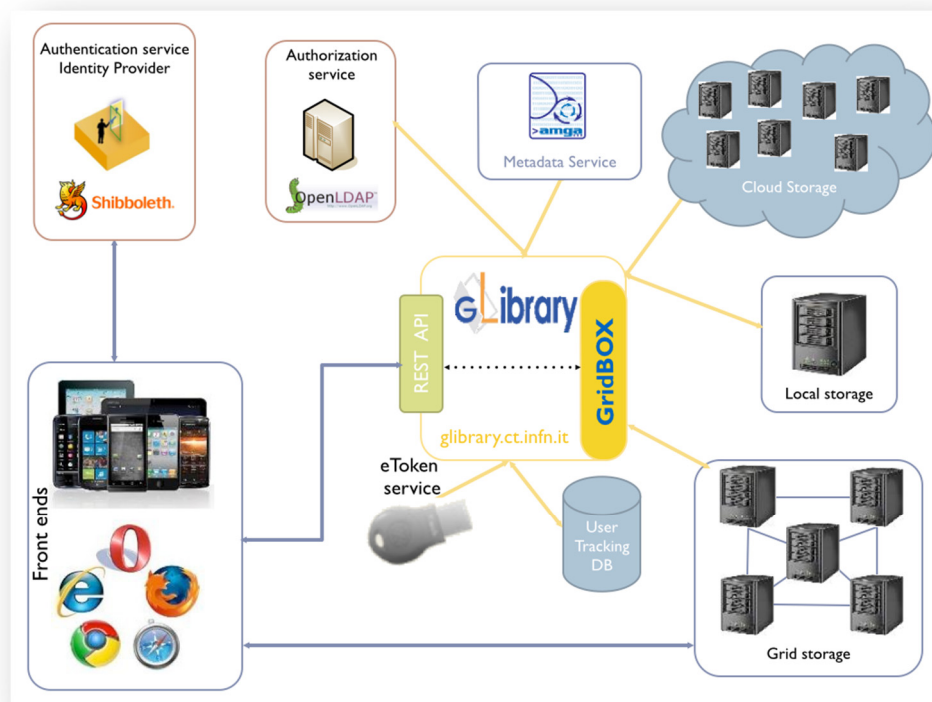


Figure 4. 6 gLibrary current architecture

¹⁸ <http://wsgi.readthedocs.org/en/latest/>

¹⁹ <http://infn.it/products/moai>

4.6 Web and mobile clients

The introduction of REST APIs to access gLibrary services brought as a consequence the birth of several clients and front-ends to interact with the system. We can group our interfaces in three categories:

1. Web Apps (SPAs/RIAs)
2. JSR-286 ²⁰Portlets
3. Native Mobile clients

Two Single Page Applications (SPA, also known as RIAs, Rich Internet Applications) belong to the first category. The first is a Repository Browser Web App, currently deployed on top of a gLibrary server. It allows to navigate all the repositories that are managed by a single instance of a gLibrary server. It offers a browsing interface with filtering capabilities, a la iTunes, metadata inspection/editing, replica download. Currently access to this front-end requires authentication via X.509 digital certificates or a demo guest access is also provided. Possible users of this interface are repository managers and data providers. This interface still uses the old PHP implementation on the backend, while we make use of the SmartClient JavaScript framework (<http://smartclient.com/product/smartclient.jsp>) on the front-end. Figure 4.4 shows this front-end in action. It can be accessed at <https://glibrary.ct.infn.it>

The second SPA is a Repository Uploader HTML5 Web App, used to upload new assets to already created repositories. While uploads are in progress, metadata can be added to each asset, using the metadata schema of its belonging *type*. It supports the upload of multiple files at once, a useful feature to upload a large set of files. The target user is the data provider and/or the repository administrator, that have the rights to add new content to repositories. Here we experiment with some bleeding-edge HTML5 features, in particular XHR²¹ and JavaScript File APIs²². The first allows to bypass CORS restrictions (Cross Origin Resource Sharing²³), but most importantly it brings support to the PUT upload method in the browser. Most of the current upload implementations use POST

²⁰ <http://jcp.org/aboutJava/communityprocess/final/jsr286/>

²¹ <http://www.w3.org/TR/XMLHttpRequest2/>

²² <http://dev.w3.org/2006/webapi/FileAPI/>

²³ <http://www.w3.org/TR/cors/>

upload with multipart/form-data that isn't ideal for big transfers and do not offer features like partial upload, resume and chunked transfers.

XHR2 allowed us to access directly the GridBox upload REST API via PUT. File APIs instead allowed to support multiple file selections and reading metadata files (like the size or mime type) on the client side, so that we can implement precise upload progress for example, features that are generally available only with the support of the backend. HTML5 permitted us to avoid the usage of any third party plugin like Java Applets or Flash uploader. The front-end of this interface has been developed using the Sencha ExtJS²⁴ 4 JavaScript framework and can be deployed in a separate machine than the gLibrary server. It interacts with gLibrary core services using its REST APIs. Figure 4.7 demonstrates its interface.

²⁴ <http://www.sencha.com/products/extjs/>

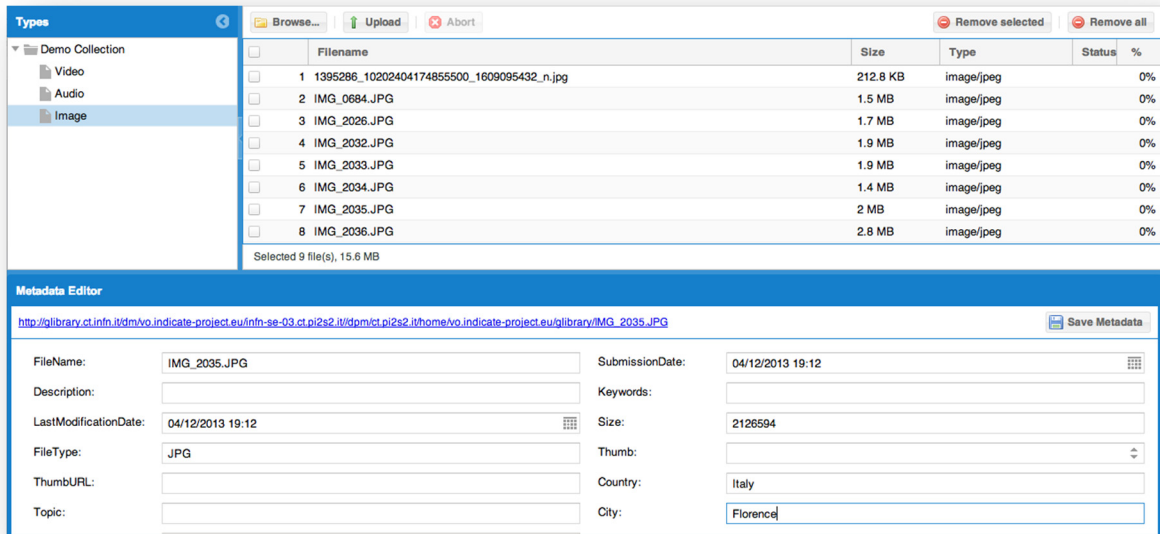


Figure 4. 7 gLibrary Uploader Web App

The second category of interfaces have been developed as a set of JSR-286 portlets, with the goal to be deployed in any Java Application Server capable of hosting portlets. Several portal software, like Liferay or JBoss Portal, are able to deploy and integrate portlets. This requirement came from the many deployments of virtual research environments called Science Gateways, based on the aforementioned portals. Portlets are like small black boxes that can be easily installed and exchanged in any portlet container with a few clicks.

Two gLibrary portlets have been currently created. The first is a Repository Browser Portlet, that offers a subset of the same features available through the Repository Browser Web App. This is targeted to Generic Authorized Users, with the need of consulting the repositories they have access to. Moreover assets can be downloaded selecting a replica from a map of storage servers. This portlet has been written in pure JavaScript using the Sencha ExtJS 3 framework and using the gLibrary REST APIs to retrieve data from the backend. Figure 4.8 shows this front-end.



Figure 4. 8 gLibrary Repository Browser portlet running in a Liferay portal

The second portlet allows to upload and edit metadata for new assets. This is actually the very same Repository Uploader HTML5 Web App wrapped into a Portlet, in such a way it can be easily deployed into any Portal.

This two portlets share the authentication and authorization system provided by the Portal where they are installed and accessed.

Finally we have developed native mobile clients both for iOS and Android platform, that permit to access gLibrary repositories on the go, with filtered browsing, metadata inspection, replica download on the local device storage for offline access. One of the benefits of accessing assets stored in a distributed infrastructure from devices with GPS

capabilities is the possibility to automatically choose the best replica on the base of the physical distance between the user and the servers. Moreover, the built-in push notification services of those platforms provide a way to subscribe and notify users that are interested in change of metadata or availability of new assets. We have a custom iPad gLibrary browser app, developed as a use case for the community of cultural heritage researchers, that allows to annotate digitalization of manuscripts. Users can subscribe to interesting pages, and once an annotation or comment is added, they gets notified.

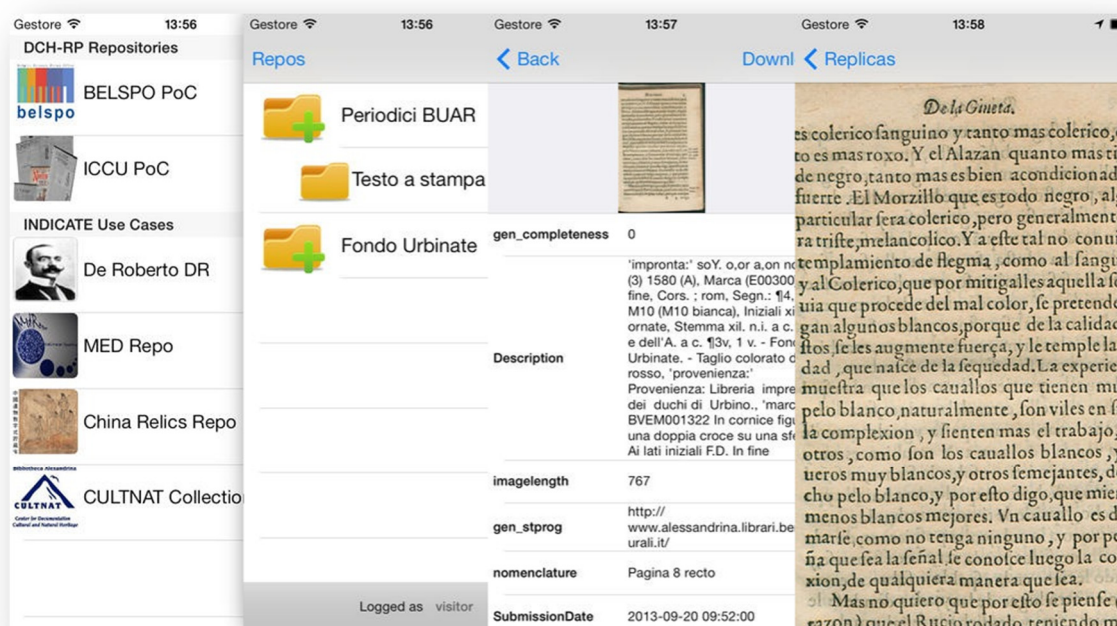


Figure 4. 9 gLibrary custom iOS app for DCH-RP project (Digital Cultural Heritage Roadmap for Preservation) repositories

This two clients are available as template apps (one for iOS and one for Android), to easily generate and publish on the App stores and Google Play market customized apps for specific communities, that want to share only a set of (or a single) gLibrary repositories. Using a simple configuration file, the app developer defines which repositories will be available for browsing and download to authorized users.

gLibrary mobile clients implement the Federeted authentication system based on Shibboleth, described in the GridBox chapter.

The iOS client for iPhone and iPad has beed built using the Appcelerator Titanium platform. Titanium is an open source project that leverages Web tecnologies such as

JavaScript, HTML and CSS to build cross-platform native applications. At the moment the target platform are iOS, Android, Blackberry 10, Tizen, HTML5, but soon the support for Windows Phone will be added. We could easily create from the same code base the Android gLibrary client. But actually the gLibrary Android client has been developed using Java and the native Android platform. Both clients communicate with the gLibrary server and GridBox using their respective APIs and exchanging data using JSON format.

Figure 4.9 and 4.10 shows the gLibrary iOS and Android browser apps, customized for the Cultural Heritage and EarthScience communities.



Figure 4. 10 gLibrary Android client for DCH-RP and EarthServer projects' repositories

Chapter 5

USE CASES

In this chapter two use cases are described, demonstrating how gLibrary and GridBox services allow communities not belonging to the IT world the exploitation of data Grids for their cataloguing and data need purposes.

5.1 Data Grids for Cultural Inheritance

Traditionally, the main goal of libraries, archives, and museums has been to preserve and spread the knowledge and the cultural heritage. Today, this really important function can take advantage of the development of digital technologies. In fact, digitization allows avoiding deterioration phenomena of original copies, caused by the handling of physical material, creating an accurate reproduction of the original ones, ensuring the usability of the contents and improving the access to documents such as books, manuscripts, parchments, leaflets, cartographies, musical scores, photos, paintings, videos, audios and three-dimensional objects.

The advantages of the digitization are manifold: access granted to a wide audience to huge collections of rare documents, preservation of original texts, reaching documents physically faraway. The prime focus of digitization is to ensure that the original media are preserved. The main interest is keeping in electronic (digital) format the documents in their integrity and in a reliable way, so to guarantee a long-term preservation, saving the original media. Finally the digitized media are themselves the addressee of secure keeping and preservation techniques, to ensure longevity and availability in time.

Converting information and data, belonging to a certain cultural heritage, into a digital format, plays an important role in communication and didactic activities, making data available on a larger scale. We can say that digitization would represent a primary tool to

allow a wide and distributed access to culture, being the only resource we have to provide next generations with such pieces of information.

Data Grids offer redundant and huge distributed storage capabilities, providing an ideal and secure place for the long-term preservation of digitized literary works and documents of artistic and historical relevance. Grid authentication and authorization mechanisms allow a fine-grained access control to archives by single users, groups or entire communities. Moreover, metadata services allow for a structured organization of digitizations for quick searches.

5.1.1 Digitization and data preservation

The amount of data produced out of the digitization process varies depending upon the number and type of documents. Several scanning workstations, working in parallel, are able to create up to several hundred of Gigabytes per working day (considering digital copies of manuscripts taken with very high resolution and colour depth).

This process creates a remarkable amount of data to store in a secure and reliable place. To meet these objectives, it is needed to have more than one copy of the digitized document (redundant storage), creating digital clones.

Due to the very high speed with which technology becomes obsolete, a particular care has to be put on the storage strategies and backup methods, to guarantee the access to the data to the next generations. Innovation and technology, in fact, enables better performances and high-level computation, but on the other hand, the risk of handling formats becoming obsolete (and not usable), is always around the corner.

Finally, the search engines developed for the users (made possible thanks to the integration of files and metadata containing information about the document, plus keywords), would allow an easy access to data, even for non-specialists.

A use case has been considered to demonstrate how Grid digital libraries can guarantee enduring preservation of literary heritage: the archives of the work of Italian writer Federico De Roberto (1861-1927) [Dig98, Cas10, Gan05], made up of almost 8000 scans, hosted on the Sicilian e-infrastructure of the COMETA consortium.

5.1.2 *Federico De Roberto works*

Among the several different components of the cultural heritage to preserve through the digitization, manuscripts surely represent a fundamental element.

The ancient hand-written papers constitute maybe the most valuable documentation of a particular author and its time. The manuscripts are single-exemplary (no copies or duplicates generally available) and they are the basis for many researches in literature. Making available historical and philological aspects of ancient books through the conservation and digitization of such documents, with the support of cutting edge technology, surely guarantees the possibility to study, analyze, and compare artists from a wider perspective, keeping data for long time.

In this particular case, the "Fondo De Roberto", a manuscript collection physically stored in the library of the "Società di Storia Patria per la Sicilia Orientale" (Catania), contains documents to be acquired and preserved such as manuscripts, rough drafts, notes, newspaper articles, letters, geographic maps, press cuttings, photographs, summing up to approximately 10,000 documents.

The size of most part of the documents is between A5 and A3, with some exceptions (7x7 cm, 3x15cm etc). Many additional sheets have been added to some documents, and there are pages with lateral notes glued to margins and refolded. Finally the digitized documents had sometimes hand corrections, cuts and notes on them. When a document was completely digitized and stored (including collected cuttings, notes, additions and complements), a set of metadata was attached to it, for future reference and search purposes. Beside metadata automatically extracted by scanning devices (ImageWidth, ImageHeight, XResolution, FileSize, CreationDate and ModifyDate), additional metadata have been inserted according the XMP²⁵ format (eXtensible Metadata Platform), a XML management format used by Adobe, compliant with the Dublin Core [Dub08] standard. The insertion has been performed in batch with the aid of Adobe Photoshop CS. The set of metadata includes Document, Title, Author, Description, Description Writer, Keywords, Copyright Status and

²⁵ XMP: <http://www.adobe.com/products/xmp/pdfs/whitepaper.pdf>

Copyright Notice. The exported images have been processed to create PDF files, optimizing them for an online browsing.

5.1.3 Implementation of the De Roberto digital repository with gLibrary

Once the acquisition stages have been concluded, there was the problem of finding a way to make available online the scanned images, in order to ensure their access according a predefined utilization strategy. The main requirements were of having a system capable of ensuring continuous availability on the net and providing the consultation from all over the world through simple and intuitive interfaces. Those requirements have naturally led to the choice of using a Grid infrastructure and gLibrary for storage and handling purposes.

The flexibility and the customization features offered by those type of systems made it possible to achieve the archiving and cataloguing on the net of the De Roberto valuable manuscript collection. Adopting the Grid resource offered by the COMETA consortium [15], the De Roberto digital repository (DRdr) was created. A beta version of the front-end is available at <https://glibrary.ct.infn.it/>.

This repository was implemented using the first release of gLibrary, that makes still use of personal X.509 certificates for authentication, and Java Applets to handle assets upload and download.

Through the gLibrary Uploader Web front-end both the high quality scans and the lower-resolution PDFs were uploaded on the Storage Elements of the COMETA infrastructure. At the same time, all the metadata associated with each image were stored on the Grid, in the AMGA Metadata Server.

Once a user successfully logs in using his personal X.509 certificate, he is able to interact with the browsing interface. It offers an easy search system among the assets, represented by the digitized manuscripts: the contents are organized by physical and semantics characteristics through the definition of types and categories and results are filtered by the selection of one or more attributes from given lists. Example of filters are: DocumentGenre, the work genre (i.e.: short story, essay, lyric tragedy); Title, refers to the

work title; FileType, scans file format (tiff, pdf, jpg, etc.); ScanQuality, gives information on the quality of the scan; DocumentType, specifies the kind of the document (manuscripts, typescripts, draft, etc.); PublicationStatus, PublicationYear, Publisher, reports the publication status (published/unpublished) the publication year and the name of the publisher; Location, provides the physical location of the documents, the number of the envelope within which it is contained.

The selection of filters, with cascading application, allows to find easily a document: the choice of a filter value, dynamically influence the values of the adjacent filters, narrowing at the same time the result set, and locating quickly the desired item.

It follows two examples to illustrate how filtering can be applied to quickly find a needed document or a group of them.

First case. A scholar may need to look for all the draft printed in the 1919. He will select the type “Scanned documents” from the Type tree, then the subtype “Printed Draft”. Choosing the PublicationYear as filter, and selecting 1919 as values among the available years, it will give a result set of all the assets satisfying his request. The search can be furtherly-refined choosing Publisher as second filter, to group drafts by Publisher.

Second case. A linguist may need to undertake the study and the analysis of “varianti d'autore” of the Ermanno Reali text. She will need to find out all the available existing versions, from the manuscripts to the published work.

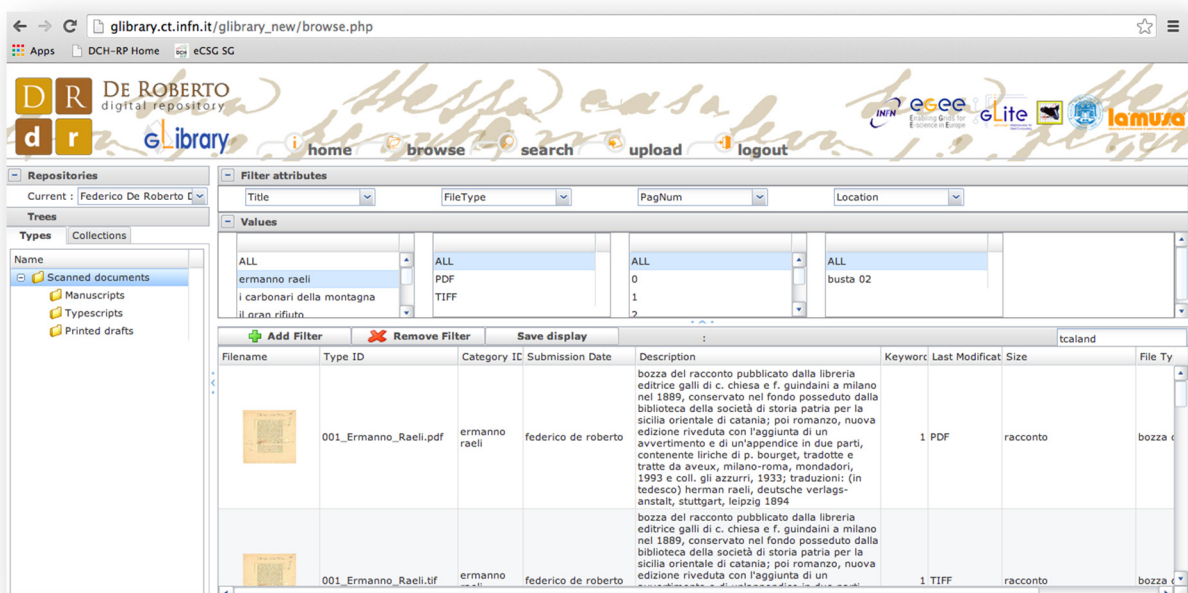


Figure 5. 1 An example of browsing using filtering to locate the work “Ermanno Reali”

Through the selection of “Scanned documents” among the types, and setting Title as filter with the value Ermanno Reali, she will be able to retrieve the list of manuscripts and all the versions of printed draft, with related modifications and corrections if any, of selected work, ordered by year

With a browsing system like this, it is possible to make search on contents while taking into account at the same time the physical preservation of documents, scan quality, physical format, etc.

Each item of the result set can be inspected with a single click. A pop up window will show all the metadata bound with each single scans, the ones added at upload time with the addition of file system and authorization metadata.

Once the document requested has been located among the results, the download can be started immediately with a click on one of the available replica links. A Java applet will retrieve the file from the right Storage Element using the certificate proxy of the user for authentication and authorization purpose. Archiving on Grid Storage Elements ensure data safe-keeping and high- availability, through the usage of replicas spread over multiple servers in different geographical locations, and secure data preservation, achieved by the high reliability of the storage systems through the usage of redundant disks and by the authentication and authorization system from the underlying Grid technologies.

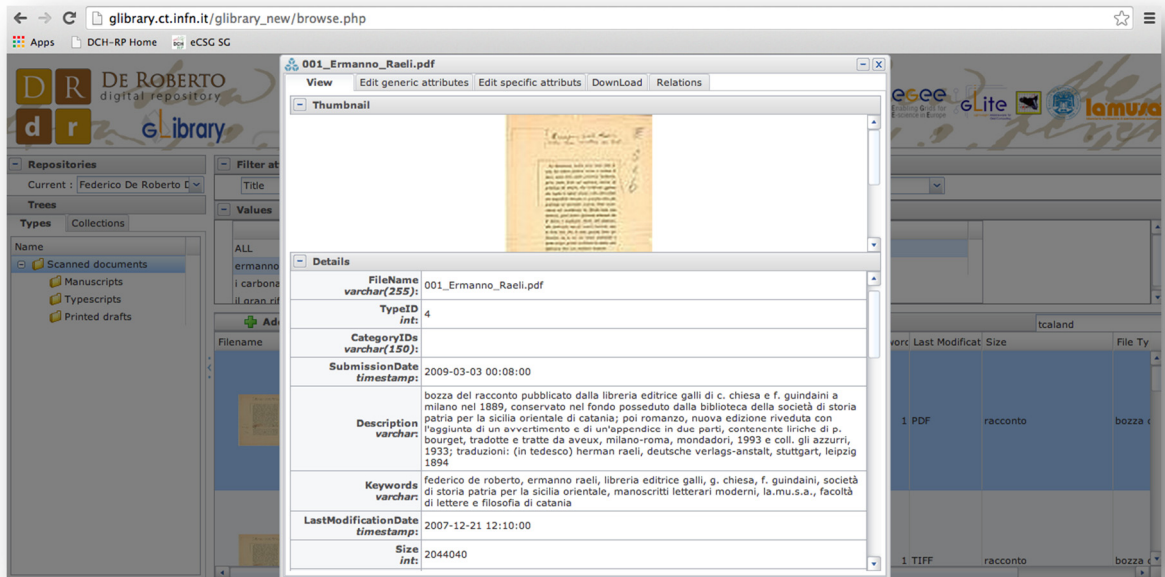


Figure 5. 2 Metadata detail view of the first page of “Ermanno Raeli” manuscript

With the implementation and the deployment of the De Roberto digital repository on a Grid infrastructure through the gLibrary platform, we have demonstrated how human science can take advantage of new technologies to achieve long term preservation of data and offer an always-on availability of contents to a huge audience widespread around the globe, that doesn't need almost any technical knowledge of the underlying technologies to use the system.

5.2 Data Grids for e-health

The field of medical imaging has developed enormously in the past 20 years. Image databases, made of thousands of medical images, are currently available to be used as a reference for individual diagnosis. At the same time, sophisticated and computationally intensive algorithms have been developed to extract information, invisible to the naked eye, from medical images. In particular, brain diseases are good candidates to benefit from such applications. Highly prevalent and burdensome chronic conditions, such as Alzheimer Disease (AD) and other neurodegenerative and neurodevelopmental disorders, can be early diagnosed by means of image- and signal-based markers of structural and functional brain changes, allowing early pharmacological or rehabilitative interventions.

Unfortunately, neuroinformatics advancements require high computational and storage resources as well as large reference image datasets of normal patients, limiting its spread to advanced academic hospitals and research centres equipped with appropriate human expertise and hardware facilities.

The aim of the Diagnostic Enhancement of Confidence by an International Distributed Environment (DECIDE) project [ABC+12], co-funded by the European Union under its Seventh Framework Program, is to design, implement, and validate a dedicated e-Infrastructure based on the Pan-European backbone GÉANT and the National Research and Education Networks (NRENs) and relying on the European Grid Infrastructure (EGI) and the National Grid Initiatives (NGIs).

Over this e-Infrastructure, a production quality service is provided around the clock for the computer-aided extraction of diagnostic disease markers for AD and schizophrenia. DECIDE offers access to a big distributed reference databases (850 and 2,200 datasets of normal and neurological subjects, respectively), large distributed computing and storage resources (more than 1,000 CPU cores and 70 TB of disk storage), and intensive image and signals processing tools.

DECIDE is focused on supporting neurologists and physicians involved in the assessment of neurodegenerative diseases in the diagnosis and prognosis and aims at enhancing user confidence by improving the reliability of the required analysis and by integrating different clinical approaches. It has been conceived to target a non-technical medical audience and aims to support the daily needs of neurologists while dealing with their patients, going well beyond the world of research.

Four applications have been deployed on the DECIDE infrastructure, with the goal to provide doctors at peripheral hospitals with service tools for determining clinical markers for the early diagnosis of neurological and psychiatric disorders (neurodegenerative diseases and schizophrenia) together with their prognostic relevance:

- **GridSPM** [CCS+09]: specifically designed for SPECT and PET neurological clinical images, provides a statistical analysis on a single-subject, based on Statistical Parametric Mapping (SPM) for the early diagnosis of Alzheimer Disease and other neurodegenerative diseases;

- **GridANN4ND** [TAA+06, BTH+08]: concerns the analysis of PET biomarkers in Neurological and Psychiatric Disorders and provides a single-subject classification of suspected patients through the use of an Artificial Neural Network;
- **GridMRIseg** [MTA+08]: implements an automatic algorithm for the subcortical segmentation of single-subject MRI brain images for hippocampal volume estimation, using the auto context model (ACMAboost) developed by LONI ;
- **GridEEG** [BCC+01, BFB+09, BKK+10]: based on a comparison of pathological versus normal subjects, implements EEG processing algorithms with the aim of detecting early symptoms of AD and distinguishing different forms of degenerative impairment.

Furthermore, the project aims to design and implement a multimodal repository, to include MRI, PET/SPECT and EEG datasets and make them available for exploitation by the data analysis software of the diagnostic/prognostic services. Medical data ownership remains in control of the physicians who contribute with their medical data to the medical repository, uploading data and reports according to their relevant authorization rights. By design, no free download of medical data from the DECIDE repository can be possible. Experts and scientists can only use the medical data contained inside the repository through the DECIDE diagnostic/prognostic services according to their authorisations.

DECIDE applications and tools are exposed to the end users (neurologists, physicians, and scientists in general) through a Science Gateway

Besides the configuration of applications, attention had therefore to be paid to the roles that users can play in accessing the DECIDE services, as well as to the authorization and permission policies enacted by the DECIDE portal. Four different kinds of users have been identified: neurologists, physicians, scientists and data managers each having different roles in the analysis process.

In particular, Data Managers can populate and update the DECIDE repositories, maintaining data and metadata, and train and update the artificial neural networks used by some applications.

A number of issues had to be tackled and solved in order to define a general architecture of the deployment of the DECIDE services on the Grid infrastructure. Two different types of services are available through the DECIDE portal:

- Front-end services, devoted to clinical and scientific users for running diagnostic services supporting neurologists in their diagnoses. Three front-end services are provided: the management of diagnostic sessions (used by neurologists), a basic application execution (used by physicians), and an advanced application execution (used by scientists).
- Back-end services, designed to be used by data managers for managing and maintaining reference databases, and for training artificial neural networks. Data managers for each application have different implementations of the same back-end service.

The DECIDE back-end services are designed mainly with the aim of providing tools for the maintenance of data and algorithms used by applications. Through these services, back-end users (data managers and developers) can either access reference databases or upgrade algorithms related to different applications.

The data manager user is provided with tools for adding and removing data concerning reference subjects and for updating the related metadata. The user can access only the specific data in the reference database which he/she owns and is responsible for. Data in the databases are shared just for statistical purposes and can be accessed by the applications during their execution on the Grid worker nodes.

Data managers can also run or update the training of the artificial neural networks used in GridANN4ND and in GridMRIseg. They can of course update only the specific application(s) they responsible for.

5.2.1 The DECIDE Science Gateway

As anticipated in the previous section, DECIDE services are exposed to users through a Science Gateway. According to [WGK+08, Wil07], a Science Gateway is a “community-development set of tools, applications, and data that is integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community”. Science Gateways are playing an important role in scientific research performed using e-Infrastructures and their relevance will further increase with the development of more sophisticated user interfaces and easier access mechanisms. Through the highly collaborative environment of a Science Gateway, users spread around the world and belonging to various Virtual Research Communities can easily cooperate to reach common goals and exploit all the resources of the cyber-infrastructure they are entitled to use.

The DECIDE Science Gateway is built within the Liferay web portal framework and portlet container and it is fully compliant with the JSR 286 (“portlet 2.0”) standard. Liferay is currently the most used framework to build Science Gateways in the “Grid world” and ships with more than sixty portlets that can be easily combined (mashed-up) to build complex and appealing e-collaboration environments.

5.2.2 Authentication and Authorisation of the DECIDE SG with Shibboleth and robot certificates

The most important requirement of the DECIDE Science Gateway was to ease the access to the distributed computing and storage resources by the largest possible community of (non IT- expert) clinicians through a set of well-defined and domain specific applications. In order to meet this requirement, authentication and authorisation mechanisms have been conceived to provide a smooth access to the applications yet preserving the security level requested by the distributed e-Infrastructure and the typology of the sensible information (clinical data) managed. Indeed, the neurological data stored in the Science Gateway have

extra requirements in terms of security, anonymity and confidentiality. It must always be clearly defined who can access which images for his/her own analysis.

Therefore, several web and Grid technologies have been adopted and deployed to ensure that the authentication and authorisation mechanisms fulfill the stringent requirements and implement the expected roles and the corresponding privileges. Such tools implement an authentication/authorisation hierarchy moving from the web to the physical Grid resources in order for users to have to interact only with the top level tool, this taking care of transferring the credentials to lower levels.

The highest component in the authorisation/authentication hierarchy has to be integrated in the Science Gateway and has to support a Single Sign On (SSO) mechanism across all services a given user is entitled (i.e., has the right) to use, in order not to confuse non-experienced users with different sets of credentials.

Many web tools support SSO within a centralised or distributed authentication framework. Nevertheless, in order to comply with currently adopted standards and support the most relevant Identity Federations in Education and Research, the DECIDE Science Gateway is compliant with the Security Assertion Markup Language (SAML) OASIS standard for credentials communication. The Shibboleth implementation of SAML has been adopted in the DECIDE Science Gateway and a library has been developed to make Liferay manage the Shibboleth token.

In many countries there is currently a big effort to create Identity Federations gathering all education and research institutions to simplify and ease the access to services for users working in different locations. Actually, they are generally managed by National Research and Education Networks in EU countries and aim at the integration of networks, services and users. Therefore, it was important for DECIDE to follow this trend in order to allow its integration with other services and increase the number of potential users.

The use of Shibboleth allows an easy integration with Identity Federations (IdFs) and individual institutions (i.e., Identity Providers – IdPs) wishing to include the DECIDE Science Gateway as one of the resources (i.e., Service Providers - SPs) for their users. When a user tries to use one of the DECIDE applications available on the Science Gateway, he/she is redirected to a Discovery Service (DS) listing all the supported IdFs and Identity

Providers (IdPs) among which he/she can select the one he/she is member of. The IdP identifies the user, generally through a pair of username and password. If the authentication by the IdP is successful, the control is returned to the Science Gateway where user authorization is checked.

The DECIDE Science Gateway is a Service Provider of the GridP (Grid IDentity Pool), a “catch-all” federation operated by INFN Catania and Consorzio COMETA to manage several Science Gateways and generic web portals. This is the federation collecting the IdPs of institutions which are not members of any Identity Federation and currently includes INFN (the Italian Institute for Nuclear Physics) and maat-g (a private company partner of the DECIDE consortium) IdPs. Besides GrIPD, the DECIDE Science Gateway is a SP of both the Italian official identity federation IDEM, one of the biggest Shibboleth-based federations available, and of the eduGAIN²⁶ interfederation gathering the many-million members of the IdFs of Belgium, Brazil, Croatia, Czech Republic, Finland, Germany, Greece, Hungary, Italy, Norway, Spain, Sweden, Switzerland, and The Netherlands.

Once a user is authenticated, the authorisation system verifies his/her credentials. The Scientific Board of DECIDE grants authorisations and a centralised LDAP-based registry, connected to the DECIDE Science Gateway, has been created to store and manage roles and privileges. User roles are then mapped onto those performing Grid transactions using the Virtual Organisation Membership Service (VOMS) functionalities. The VOMS-based Virtual Organization vo.eu-decide.eu has been created and four Groups have been defined, one for each application. Additionally, for each group, four Roles have been defined that represent the four different kinds of users of the DECIDE services explained in the.

Besides the few developers of DECIDE applications, who need low-level access, the DECIDE VOMS is populated only by robot certificates who are used to sign Grid transactions with their proxies.

The management of personal certificates to access e-Infrastructures has demonstrated to be difficult by non-expert users and represents a limiting factor to the rapid spread of this technology in new scientific domains where computer science is not a basic knowledge. A notable step forward to make the access to Grid infrastructures as much transparent and smooth as possible has recently been achieved with the introduction of

²⁶ <http://www.geant.net/service/eduGAIN/Pages/home.aspx>

robot certificates, also referred as portal certificates, and their integration inside traditional general-purpose portals and Science Gateways. The advantages introduced by this new kind of digital certificates are manifold and ten Certification Authorities in Europe have already adopted them. For security sake, robot certificates are usually stored on board of tamper-resistant devices such as smartcards. This improves the security and avoids any fraudulent use of the private keys. A multi-threaded server, called e-Token server, has been created and configured to manage a list of robot certificates (one certificate per application) stored in different USB.

eToken PRO 32/64 KB smart cards released by SafeNet²⁷. The eToken Server provides the DECIDE Science Gateway (and other gateways developed at INFN Catania) with a 24 × 7 service and holds the web services to access the smart cards and interact both with the Virtual Organisation Membership Service and the automatic proxy renewal (MyProxy) service. This crypto library developed combines different standards and programming native libraries with the latest release of robot certificates. A Java multi-platform client, configured for inter-service communication via HTTPS, completes the architecture. The e-Token server is built on top of the Apache Tomcat Application Server and configured to accept requests only from a set of authorized “clients” (i.e., the Science Gateways). This ensures scalability and high performances especially when the server has to deal with huge numbers of requests. To further improve its performances and reduce the waiting time to get a proxy, the eToken Server implements also a cache mechanism.

The usage workflow of the “light-weight” Grid crypto library used by the DECIDE Science Gateway is shown in Fig. 5.3.

²⁷ <http://www.safenet-inc.com/>

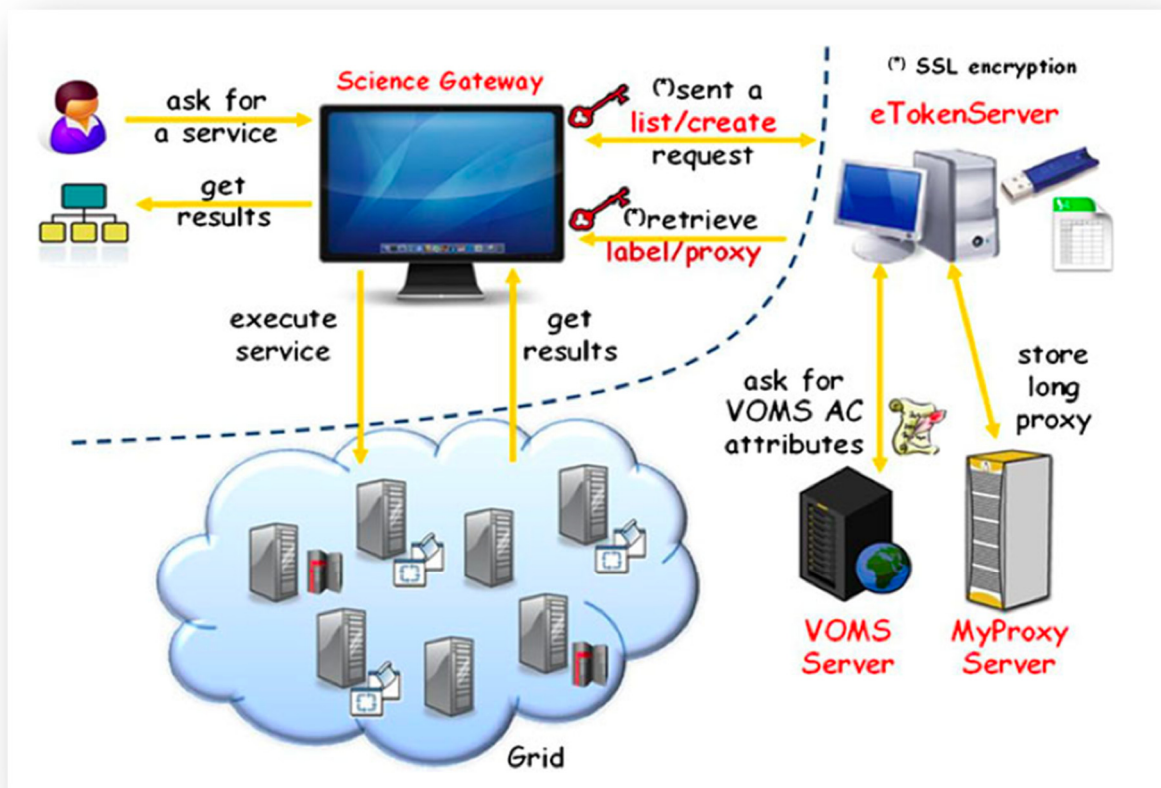


Figure 5. 3 The workflow for implementing Grid authentication by means of the eToken Server

Once the accredited IdP has successfully authenticated the user, his/her authorisation rights on the local LDAP registry are checked and, if the verification is successful, he/she is logged into the DECIDE Science Gateway. Then, according to the information stored in the LDAP registry and the application(s) he/she wants to run from within the portal, the Science Gateway sends a requestID to the eToken Server. If the Science Gateway is authorised, and taking into account the information available in the HashMap, the eToken Server sends back to the Science Gateway a new proxy certificate, if the requestID is not found or the lifetime of the old proxy is expired. Otherwise it sends a cached proxy, if a valid proxy certificate is available in memory. This operation is completely transparent from the end-user point of view, and allows the user to execute the operations requiring the usage of the Grid services and get the results. The retrieval of a cached proxy takes only 20 ms while the creation of a new one requires about 5 s.

5.2.3 *The gLibrary Repository Manager*

One of the requirements of the DECIDE project is the creation and access of “reference databases” containing structured data related to normal patients’ images. Those databases needed to be recreated in the adopted infrastructure, populated on the local LDAP registry are checked and, if the verification is successful, he/she is logged into the DECIDE Science Gateway. Then, according to the information stored in the LDAP registry and the application(s) he/she wants to run from within the portal, the Science Gateway sends a requestID to the eToken Server. If the Science Gateway is authorised, and taking into account the information available in the HashMap, the eToken Server sends back to the Science Gateway a new proxy certificate, if the requestID is not found or the lifetime of the old proxy is expired. Otherwise it sends a cached proxy, if a valid proxy certificate is available in memory. This operation is completely transparent from the end-user point of view, and allows the user to execute the operations requiring the usage of the Grid services and get the results. The retrieval of a cached proxy takes only 20 ms while the creation of a new one requires about 5 s.

One of the requirements of the DECIDE project is the creation and access of “reference databases” containing structured data related to normal patients’ images. Those databases needed to be recreated in the adopted infrastructure, populated and later queried and accessed by other applications running in the worker nodes of Grid sites. In particular, four databases have been created and deployed on the DECIDE Grid infrastructure for each of the following types of data:

- PET/SPECT;
- EEG;
- MRI;
- PET/SPECT, to serve the GridANN4ND □ application.

Image files of normal patients need to be stored on the distributed storage resources of the infrastructure in an encrypted format. The aforementioned reference databases actually represent the metadata of those binary image files. For the sake of simplicity, we prefer to use the term “digital asset” to refer to the entity made of a binary image file and its set of metadata, and so we use the term “digital repository” (or simply “repository”) to indicate a set of digital assets.

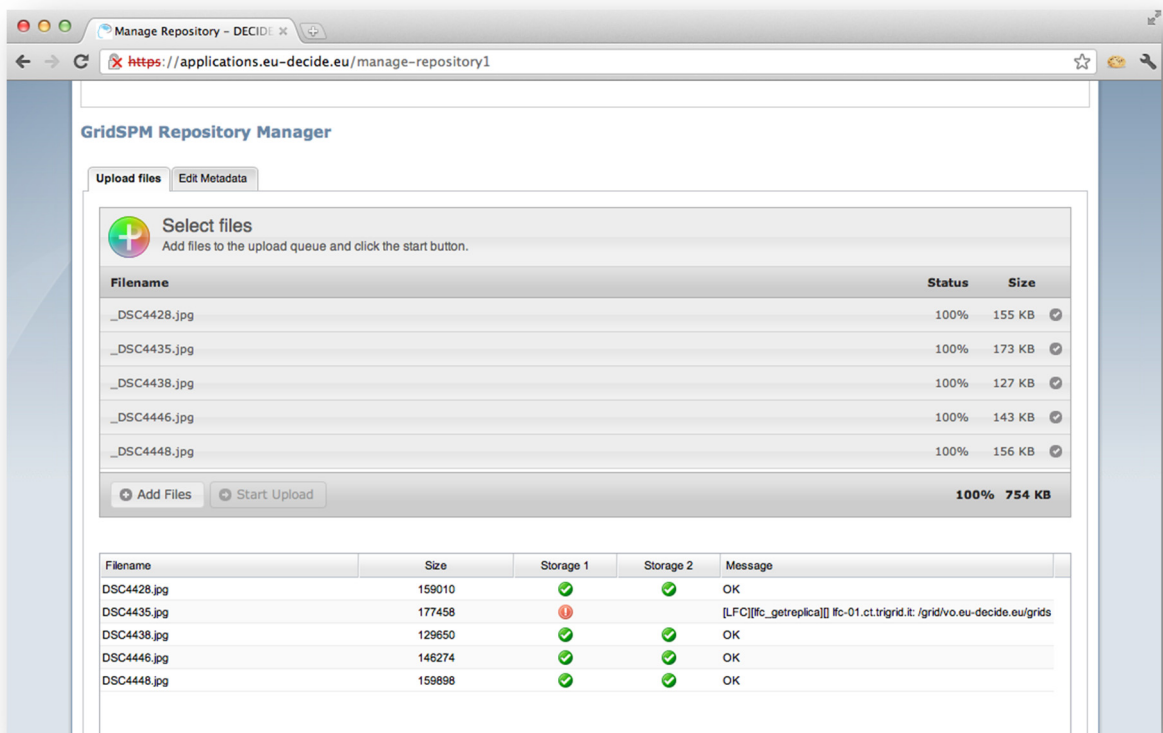


Figure 5. 4 File encryption, uploading and replication with the Repository Manager

Four repositories have been created on the DECIDE e-Infrastructure, with the same names of the four reference databases. In order to achieve this, we have used the gLibrary platform to create, access and manage digital assets on Grid. gLibrary, in fact, presents to users and applications the abstraction of digital asset as a unique entity, handling “behind the scenes” transactional interactions with the AMGA metadata catalogue for managing relational metadata and to the Grid Storage Elements for managing binary files.

Data Managers of DECIDE need to access the repositories to add assets, by uploading anonymised patient data in an encrypted format and adding metadata from the reference database to those files. As those users are supposed not to be Grid experts, and being command line access considered forbidden by requirement, we developed an easy-to-use graphical interface in the form of a set of portlets integrated in the DECIDE Science Gateway. These portlets provide access both to the Secure Storage System [SS07], used to encrypt the uploaded files and save the encryption keys into a keystore service located within the boundaries of hospitals, and to the gLibrary server that registers those files as assets and allows the addition and editing of their metadata.

The interactions between the Liferay portlets and those two services are managed through their REST-based front-ends.

Figure 5.5 Repository Manager Metadata Editor

To this goal, we have developed a set of APIs to perform all the available operations: encryption and replication of files to random Storage Elements of the infrastructure, addition and editing of their metadata.

The authentication of the APIs is done through the same Shibboleth-based mechanism used to access the DECIDE Science Gateway so a user trying to use the API without a valid Shibboleth token in his/her client (i.e., the web browser) would be refused.

For the sake of completeness, the steps and the interactions among the service to achieve the upload of an asset are explained below:

1. A user successfully logged in the Liferay-based Science Gateway retrieve the DECIDE Repository Manager portlet in his/her browser;

2. The user selects one or more files from his/her local filesystem (batch and asynchronous uploads are supported too) and transfer them via HTTPS to the “Uploader service” deployed within the boundary of the hospital he belongs too. One strong requirement of the service is that patient images should never cross the boundaries of the hospital that owns them in an unencrypted format.

3. The encrypt command is issued to the REST API of the Science Gateway. This API will act as a proxy to the machine providing the actual service, because of the Cross Origin Resource Sharing problem that denies access to resources outside of the domain from which the original web content (the portlet in our case) has been retrieved;

4. The encrypted request is forwarded back to the REST API of the “Uploader Service”, the service in charge of encryption and upload to the Grid storage resources;

5. The “Uploader Service” ask for and download a proxy certificate, according to the rights of the logged user (DataManager Role in our case) issued by the eToken Server (myproxy.ct.infn.it) needed for the subsequent Grid interactions;

6. The “Uploader Service” starts the encrypted upload, storing the encryption key into the keystore service (deployed in the same keystore.ct.infn.it machine), and moving the encrypted file to the first random storage element of the e-Infrastructure; notice that the files moved outside the hospital are transferred encrypted over the GSIFTP protocol;

7. A Logical File Name is assigned to the encrypted file and saved into the LCG File Catalog (LFC) that will keep track of this and his replicas;

8. At the end of the upload process, the client issues a “Replicate” command to the REST API of the Science Gateway;□

9. 10. 11. In a similar way, the Replication request is forwarded to the REST API of the Uploader Service that in turn replicates the current uploaded file to a second random storage element and keep note of the destination on the Logical File Catalog service.

12. Once both the encrypted upload and replication phases are completed (Fig. 4), the user can fill all the metadata related to the uploaded file. All these requests are issued against the REST interface of the Science Gateway.□

13. In turns, the Science Gateway forwards the metadata management requests to the gLibrary REST APIs.□

14. gLibrary retrieves a proxy certificate from the eTokenServer according to the authorization of the logged user□

15. On behalf of the user, the gLibrary service executes the proper sequences of commands towards the AMGA Metadata Service of the DECIDE Grid infrastructure.

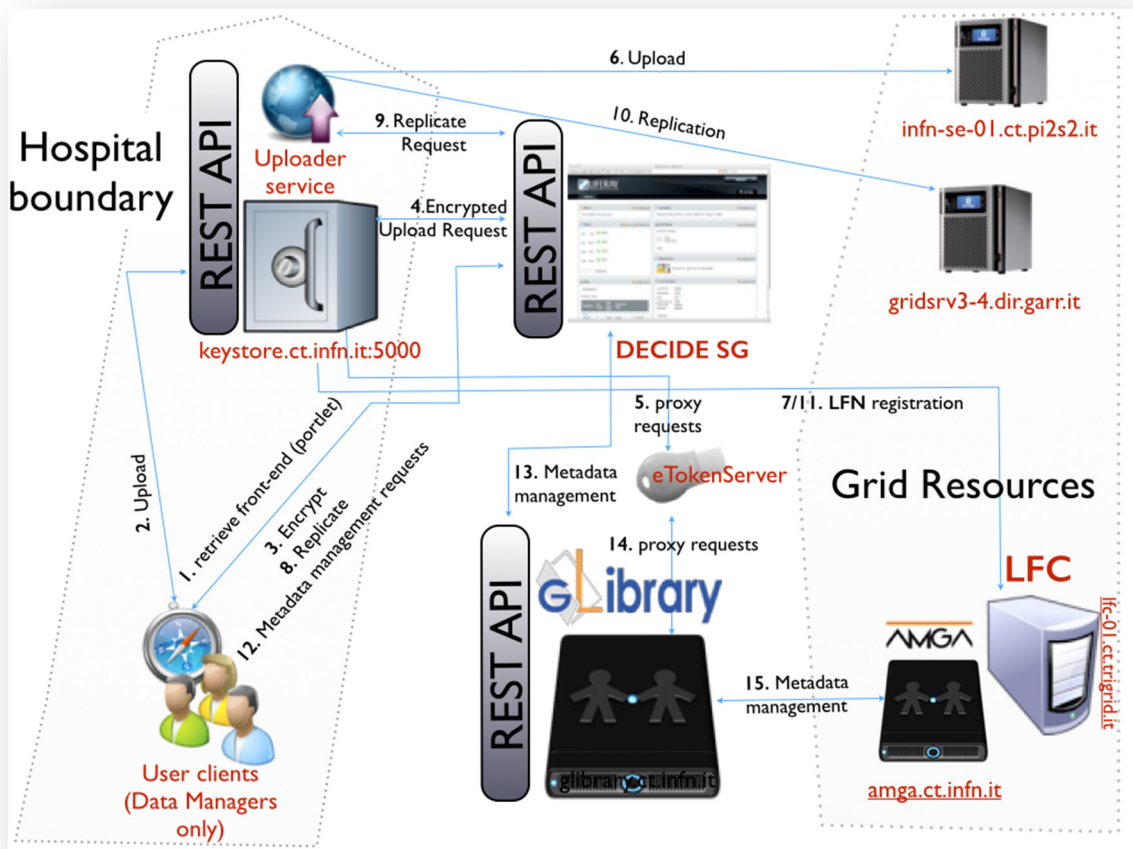


Figure 5. 6 the Respository Manager Service architecture

The main goal of the DECIDE project is to uptake and exploit the e-Infrastructure paradigm in order to provide a dedicated production quality service for computer-aided diagnosis and research in the field of neurological diseases. DECIDE builds upon GÈANT and EGI with the aim of fulfilling the specific needs of the neuroscientific and medical community. This will provide the community with new diagnostic and research tools and enable clinicians to tackle new challenges in their domain.

The services that have been realized by the DECIDE project are exposed to end users through a new type of Science Gateway based on the Liferay portlet container and on worldwide standards which makes use of a sophisticated authentication and authorization mechanism able to ease the access and use to Grid yet implementing a fine grained control on roles and corresponding privileges. The DECIDE Science Gateway allows also the

creation and management of large distributed repositories of medical images with the possibility to encrypt the stored data, through the gLibrary platform.

Chapter 6

CONCLUSIONS

In this thesis, we have demonstrated how the use of common standards can simplify the access to a complex technology such as Grids. In particular we were focused on data and metadata management services, but some work has already been done in the area of Grid computing too.

A simpler authentication system is a key element to let as many researchers, scientists, students and even casual users as possible to approach a new technology, that lets users spread around the world and belonging to different organisations easily to cooperate to reach common goals and exploit all the resources they need to accomplish their work.

Using standard protocols, such as HTTP and WebDAV for accessing Grid storage services poses the base for the implementation of storage aggregation and dynamic federations [FRD+12]. This allow to build a system that offers a unique view of the storage and metadata ensemble and the possibility of integration of other compatible resources such as those from cloud providers. Such so-called storage federations of standard protocols-based storage servers give a unique view of their content, thus promoting simplicity in accessing the data they contain and offering new possibilities for resilience and data placement strategies. The most exciting feature is the ability to access such as huge aggregated storage right from clients built-in consumer desktop environment and from mobile devices.

LIST OF FIGURES

<i>Figure 2. 1 Authentication and Authorization Schema of a Science Gateway.....</i>	<i>14</i>
<i>Figure 2. 2 Workflow of the registration procedure</i>	<i>16</i>
<i>Figure 2. 3 Workflow of the sign-in procedure</i>	<i>17</i>
<i>Figure 2. 4 Architecture of the Authentication and Authorization (AA) system in Science Gateways....</i>	<i>19</i>
<i>Figure 2. 5 eToken server architecture</i>	<i>20</i>
<i>Figure 3. 1 DPM HTTP redirection mechanism</i>	<i>24</i>
<i>Figure 3. 2 GridBox Architecture, sequence diagram</i>	<i>32</i>
<i>Figure 3. 3 GridBox native iOS client – Shibboleth authentication.....</i>	<i>38</i>
<i>Figure 3. 4 Accessing the Shibboleth protected WebDAV interface from a browser.....</i>	<i>39</i>
<i>Figure 3. 5 browsing the a Grid storage using a desktop WebDAV client</i>	<i>41</i>
<i>Figure 3. 6 using the built-in WebDAV client built-in in Mac OS to mount a Grid storage directory.....</i>	<i>42</i>
<i>Figure 3. 7 browsing the Grid storage through Gridbox as a local disk.....</i>	<i>42</i>
<i>Figure 4. 1 Example of types’ tree</i>	<i>47</i>
<i>Figure 4. 2 Example of Types and Attributes’ list.....</i>	<i>48</i>
<i>Figure 4. 3 Test repository for the EELA project – hierarchy of types on the left</i>	<i>49</i>
<i>Figure 4. 4 Test repository – Available filters for the Training type</i>	<i>50</i>
<i>Figure 4. 5 first prototype of gLibrary implementation</i>	<i>55</i>
<i>Figure 4. 6 gLibrary current architecture.....</i>	<i>56</i>
<i>Figure 4. 7 gLibrary Uploader Web App.....</i>	<i>59</i>
<i>Figure 4. 8 gLibrary Repository Browser portlet running in a Liferay portal.....</i>	<i>60</i>
<i>Figure 4. 9 gLibrary custom iOS app for DCH-RP project (Digital Cultural Heritage Roadmap for Preservation) repositories</i>	<i>61</i>
<i>Figure 4. 10 gLibrary Android client for DCH-RP and EarthServer projects’ repositories.....</i>	<i>62</i>

<i>Figure 5. 1 An example of browsing using filtering to locate the work “Ermanno Raeli”</i>	<i>67</i>
<i>Figure 5. 2 Metadata detail view of the first page of “Ermanno Raeli” manuscript</i>	<i>69</i>
<i>Figure 5. 3 The workflow for implementing Grid authentication by means of the eToken Server</i>	<i>77</i>
<i>Figure 5. 4 File encryption, uploading and replication with the Repository Manager</i>	<i>79</i>
<i>Figure 5. 5 Repository Manager Metadata Editor</i>	<i>80</i>
<i>Figure 5. 6 the Respository Manager Service architecture.....</i>	<i>83</i>

BIBLIOGRAPHY

- [AAB+11] Andronico, G., Ardizzone, V., Barbera, R., Becker, B., Bruno, R., Calanducci, A., ... & Scardaci, D. (2011). **e-Infrastructures for e-Science: a global view.** *Journal of Grid Computing*, 9(2), 155-184.
- [AAC+12] Aiftimiei, C., Aimar, A., Ceccanti, A., Cecchi, M., Di Meglio, A., Fuhrmam, P., ... & White, J. (2012, October). **Towards next generations of software for distributed infrastructures: the European Middleware Initiative.** In *E-Science (e-Science), 2012 IEEE 8th International Conference on* (pp. 1-10). IEEE.
- [ABB+03] Allcock, W., Bester, J., Bresnahan, J., Chervenak, A., Liming, L., & Tuecke, S. (2003). **GridFTP: Protocol extensions to FTP for the Grid.** *Global Grid ForumGFD-RP*, 20.
- [ABC+12] Ardizzone, V., Barbera, R., Calanducci, A., Fargetta, M., Ingrà, E., Porro, I., ... & Schenone, A. (2012). **The DECIDE science gateway.** *Journal of Grid Computing*, 10(4), 689-707.
- [ABC+12] Ardizzone, V., Bruno, R., Calanducci, A., Carrubba, C., Fargetta, M., Ingrà, E., ... & Barbera, R. (2012). **Science Gateways for Semantic-Web-Based Life Science Applications.** *HealthGrid Applications and Technologies Meet Science Gateways for Life Sciences*, 175, 119.
- [ABF+12] Alvarez, A., Beche, A., Furano, F., Hellmich, M., Keeble, O., & Rocha, R. (2012, December). **DPM: future proof storage.** In *Journal of Physics: Conference Series* (Vol. 396, No. 3, p. 032015). IOP Publishing.
- [ACC+04] Alfieri, R., Cecchini, R., Ciaschini, V., dell'Agnello, L., Frohner, A., Gianoli, A., ... & Spataro, F. (2004, January). **VOMS, an authorization**

- system for virtual organizations.** In *Grid computing* (pp. 33-40). Springer Berlin Heidelberg.
- [ACD+11] Andronico, G., Calanducci, A., De Filippo, A., De Gregorio, G., Foti, G., La Rocca, G., ... & Valente, E. (2011). **e-Infrastructures for Cultural Heritage Applications.** *Handbook of Research Technologies and Cultural Heritage: Applications and Environments*, 341-369.
- [AL99] Adams, C., & Lloyd, S. (1999). **Understanding the Public-Key Infrastructure: Concepts, Standards and Deployment Considerations.** *Sams Publishing*.
- [BB11] Bellembos, T., & Bourges, R. (2011). **The open-source ESUP-Portail WebDAV storage solution.**
- [BBL02] Baker, M., Buyya, R., & Laforenza, D. (2002). **Grids and Grid technologies for wide-area distributed computing.** *Software: Practice and Experience*, 32(15), 1437-1466.
- [BCC+01] Babiloni, F., Carducci, F., Cincotti, F., Del Gratta, C., Pizzella, V., Romani, G. L., ... & Babiloni, C. (2001). **Linear inverse source estimate of combined EEG and MEG data related to voluntary movements.** *Human brain mapping*, 14(4), 197-209.
- [BCL+05] Baud, J. P., Casey, J., Lemaitre, S., Nicholson, C., Smith, D., & Stewart, G. (2005, September). **Lcg data management: From edg to egee.** In *UK eScience All Hands Meeting Proceedings, Nottingham, UK*.
- [BDF+09] Barbera, R., Donvito, G., Falzone, A., La Rocca, G., Milanesi, L., Maggi, G. P., & Vicario, S. (2009). **The GENIUS Grid Portal and robot certificates: a new tool for e-Science.** *BMC bioinformatics*, 10(Suppl 6), S21.

- [BFB+09] Babiloni, C., Ferri, R., Binetti, G., Vecchio, F., Frisoni, G. B., Lanuzza, B., ... & Rossini, P. M. (2009). **Directionality of EEG synchronization in Alzheimer's disease subjects.** *Neurobiology of aging*, 30(1), 93-102.
- [BFR11] Barbera, R., Fargetta, M., & Rotondo, R. (2011, March). **A Simplified Access to Grid Resources by Science Gateways.** In *Proceedings of The International Symposium on Grids and Clouds, Taipei (2011), Proceedings of Science (ISGC 2011 & OGF 31)* (p. 23).
- [BHW05] Basney, J., Humphrey, M., & Welch, V. (2005). The MyProxy online credential repository. *Software: Practice and Experience*, 35(9), 801-816.
- [BKK+10] Blinowska, K., Kus, R., Kaminski, M., & Janiszewska, J. (2010). **Transmission of brain activity during cognitive task.** *Brain topography*, 23(2), 205-213.
- [BTH+08] Bose, S. K., Turkheimer, F. E., Howes, O. D., Mehta, M. A., Cunliffe, R., Stokes, P. R., & Grasby, P. M. (2008). **Classification of schizophrenic patients and healthy controls using [18F] fluorodopa PET imaging.** *Schizophrenia research*, 106(2), 148-155.
- [Cas10] Castelli, R. (2010). **Il punto su Federico De Roberto: per una storia delle opere e della critica** (Vol. 6). Bonanno.
- [CBS+09] Calanducci, A. S., Barbera, R., Sevilla Cedillo, J., De Filippo, A., Saso, M., Iannizzotto, S., ... & Vicinanza, D. (2009, May). **Data Grids for conservation of cultural inheritance.** In *Proceedings of the 1st ACM workshop on Data Grids for eScience* (pp. 1-6). ACM.
- [CCC+07] Calanducci, A. A., Cherubino, C., Ciuffo, L. N., Fargetta, M., & Scardaci, D. (2007, June). **A Digital Library Management System for Grid.** In *Enabling Technologies: Infrastructure for Collaborative Enterprises, 2007. WETICE 2007. 16th IEEE International Workshops on* (pp. 269-272). IEEE.
- [CCP+08] Calanducci, A., Castrillo, F. P., Pollán, R. R., & del Solar, M. R. (2008, September). **Enabling digital repositories on the Grid.** In *Advanced*

Engineering Computing and Applications in Sciences, 2008. ADVCOMP'08. The Second International Conference on (pp. 45-50). IEEE.

- [CCS+09] Castiglioni, I., Canesi, B., Schenone, A., Perani, D., & Gilardi, M. C. (2009). **A Grid-based SPM service (GriSPM) for SPECT and PET neurological studies.** *European journal of nuclear medicine and molecular imaging*, 36(7), 1193-1195.
- [CER12] CERN European Organization for Nuclear Research (2012, February). **CERN Certificate Policy And Certificate Practice Statement.** <https://ca.cern.ch/ca/CRL/Policy/cp-cps.pdf>
- [CF99] Chokhani, S., & Ford, W. (1999). **Internet x. 509 public key infrastructure certificate policy and certification practices framework.**
- [CFG+07] Carbone, A., Forti, A., Ghiselli, A., Lanciotti, E., Magnoni, L., Mazzucato, M., ... & Zappi, R. (2007, December). **Performance studies of the StoRM storage resource manager.** In *e-Science and Grid Computing, IEEE International Conference on* (pp. 423-430). IEEE.
- [CFK+00] Chervenak, A., Foster, I., Kesselman, C., Salisbury, C., & Tuecke, S. (2000). **The data Grid: Towards an architecture for the distributed management and analysis of large scientific datasets.** *Journal of network and computer applications*, 23(3), 187-200.
- [CJK04] Cornwall, L. A., Jensen, J., Kelsey, D. P., Frohner, Á., Kouřil, D., Bonnassieux, F., ... & McNab, A. (2004). **Authentication and authorization mechanisms for multi-domain Grid environments.** *Journal of Grid Computing*, 2(4), 301-311.
- [CLS05] Curbera, F., Leymann, F., Storey, T., Ferguson, D., & Weerawarana, S. (2005). **Web services platform architecture: SOAP, WSDL, WS-policy, WS-addressing, WS-BPEL, WS-reliable messaging and more.** *Englewood Cliffs: Prentice Hall PTR.*

- [DEF+05] Dorigo, A., Elmer, P., Furano, F., & Hanushevsky, A. (2005). **XROOTD-A Highly scalable architecture for data access**. *WSEAS Transactions on Computers*, 1(4.3).
- [Dig98] Di Grado, A. (1998). **La vita, le carte, i turbamenti di Federico De Roberto, gentiluomo** (Vol. 7). Fondazione Verga.
- [Dub08] Dublin Core Metadata Initiative. (2008). **Dublin core metadata element set**, version 1.1.
- [EC13] European Commission TERENA (2013), **A Study on Authentication and Authorisation Platforms For Scientific Resources in Europe**. <https://confluence.terena.org/download/attachments/30474266/2012-AAA-Study-report-final.pdf?version=1&modificationDate=1355503760046&api=v2>
- [EGI12] EGI VO Portal Policy: <https://documents.egi.eu/public/ShowDocument?docid=80>. Accessed Oct 2012
- [EGI12a] EGI Grid Security Traceability and Logging Policy: <https://documents.egi.eu/public/ShowDocument?docid=81>. Accessed Oct 2012
- [EKK+03] Ellert, M., Konstantinov, A., Kónya, B., Smirnova, O., & Wäänänen, A. (2003). **The NorduGrid project: Using Globus toolkit for building Grid infrastructure**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 502(2), 407-410.
- [EMI12] EMI Project (2012). **The European Middleware Initiative to support Security Token Services**. http://www.eu-emi.eu/further-readings/-/asset_publisher/7AeX/content/security-token-services
- [Eur10] The European Policy Management Authority for Grid Authentication (2010, February). **Guideline on Approved Robots**. <https://www.euGridpma.org/guidelines/robot/approved-robots-20100119.pdf>

- [FBS06] Filipovic, B., & Straub, T. (2006). **Grid Security Infrastructure**—ein Überblick. *DFN Tagungsband*, 115-126.
- [FGM+99] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., & Berners-Lee, T. (1999). **Hypertext transfer protocol—HTTP/1.1**.
- [Fie00] Fielding, R. (2000). Representational state transfer. *Architectural Styles and the Design of Network-based Software Architecture*, 76-85.
- [FK03] Foster, I., & Kesselman, C. (Eds.). (2003). **The Grid 2: Blueprint for a new computing infrastructure**. *Access Online via Elsevier*.
- [FKT01] Foster, I., Kesselman, C., & Tuecke, S. (2001). **The anatomy of the Grid: Enabling scalable virtual organizations**. *International journal of high performance computing applications*, 15(3), 200-222.
- [Fos02] Foster, I. (2002). **The Grid: a new infrastructure for 21st century science**. *Phys. Today*, 55(ANL/MCS/JA-42173).
- [FRD+12] Furano, F., da Rocha, R. B., Devresse, A., Keeble, O., Ayllón, A. Á., & Fuhrmann, P. (2012, December). **Dynamic federations: storage aggregation using open tools and protocols**. In *Journal of Physics: Conference Series* (Vol. 396, No. 3, p. 032042). IOP Publishing.
- [Fuh04] Fuhrmann, P. (2004, April). **dCache, the Commodity Cache**. In *MSST* (pp. 171-175).
- [Gan05] Ganeri, M. (2005). **L'Europa in Sicilia: saggi su Federico De Roberto**. Le Monnier università.
- [Geb05] Gebel, G. (2005). **Federated Identity: A Progress Report**. In *Isse 2005- Securing Electronic Business Processes: Highlights of the Information Security Solutions Europe 2005 Conference* (p. 3). Springer.
- [GJG+05] Gagliardi, F., Jones, B., Grey, F., Bégin, M. E., & Heikkurinen, M. (2005). **Building an infrastructure for scientific Grid computing: status and**

- goals of the EGEE project.** *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 363(1833), 1729-1742.
- [GWF+99] Goland, Y., Whitehead, E., Faizi, A., Carter, S., & Jensen, D. (1999). **RFC 2518—HTTP Extensions for Distributed Authoring—WEBDAV.** *Internet Engineering Task Force, Request for Comments Report, 2518*, 1-94.
- [HH11] Hammer-Lahav, D. E., & Hardt, D. **The OAuth2.0 Authorization Protocol.** 2011. IETF Internet Draft.
- [KLO10] Kranzlmüller, D., de Lucas, J. M., & Öster, P. (2010). **The European Grid Initiative (EGI).** In *Remote Instrumentation and Virtual Laboratories* (pp. 61-66). Springer US.
- [KSP08] Koblitz, B., Santos, N., & Pose, V. (2008). **The AMGA metadata service.** *Journal of Grid Computing*, 6(1), 61-76.
- [KT01] Kalmady, R., & Tierney, B. (2001, March). **A Comparison of GSIFTP and RFIO on a WAN.** In *Proceedings of Computers in High Energy Physics*.
- [LBC+11] La Rocca, G., Barbera, R., Ciaschini, V., Falzone, A., & Monforte, S. (2011). **A new “lightweight” Crypto Library for supporting a new Advanced Grid Authentication Process with Smart Cards.** *Proceedings of Science (ISGC 2011 & OGF 31)*, 29.
- [LEP+06] Laure, E., Edlund, A., Pacini, F., Buncic, P., Barroso, M., Di Meglio, A., ... & Fisher, S. M. (2006). **Programming the Grid with gLite** (No. EGEE-TR-2006-001).
- [MCC+04] Morgan, R. L., Cantor, S., Carmody, S., Hoehn, W., & Klingenstein, K. (2004). **Federated Security: The Shibboleth Approach.** *Educause Quarterly*, 27(4), 12-17.
- [MKR+07] Matyska, L., Křenek, A., Ruda, M., Sitera, J., Kouřil, D., Voců, M., ... & Salvat, Z. (2007). Job tracking on a Grid—the Logging and Bookkeeping

and Job Provenance services. Available on: <http://home.zcu.cz/~bonik/papers/lbjp.pdf>.

- [Mol12] Molnár, Z. (2012). **Next generation WLCG File Transfer Service (FTS)**.
- [MTA+08] Morra, J. H., Tu, Z., Apostolova, L. G., Green, A. E., Avedissian, C., Madsen, S. K., ... & Thompson, P. M. (2008). **Validation of a fully automated 3D hippocampal segmentation method using subjects with Alzheimer's disease mild cognitive impairment, and elderly controls**. *Neuroimage*, 43(1), 59-68.
- [NFS89] Nowicki, B. (1989). **Nfs: Network file system protocol specification**.
- [RHP+08] Ragouzis, N., Hughes, J., Philpott, R., Maler, E., Madsen, P., & Scavo, T. (2008). **Security assertion markup language (saml) v2.0 technical overview**. *OASIS Committee Draft*, 2.
- [RR06] Recordon, D., & Reed, D. (2006, November). **OpenID 2.0: a platform for user-centric identity management**. In *Proceedings of the second ACM workshop on Digital identity management* (pp. 11-16). ACM.
- [RRX11] Rieger, S., Richter, H., & Xiang, Y. (2011, September). **Introducing Federated WebDAV Access to Cloud Storage Providers**. In *CLOUD COMPUTING 2011, The Second International Conference on Cloud Computing, GRIDS, and Virtualization* (pp. 46-51).
- [RSA78] Rivest, R. L., Shamir, A., & Adleman, L. (1978). **A method for obtaining digital signatures and public-key cryptosystems**. *Communications of the ACM*, 21(2), 120-126.
- [SK06a] Santos, N., & Koblitz, B. (2006). **Distributed metadata with the AMGA metadata catalog**. *arXiv preprint cs/0604071*.

- [SK06b] Santos, N., & Koblitz, B. (2006). **Metadata services on the Grid**. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 559(1), 53-56.
- [SS07] Scardaci, D., & Scuderi, G. (2007, August). **A secure storage service for the glite middleware**. In *Information Assurance and Security, 2007. IAS 2007. Third International Symposium on* (pp. 261-266). IEEE.
- [SSG02] Shoshani, A., Sim, A., & Gu, J. (2002, April). **Storage resource managers: Middleware components for Grid storage**. In *NASA Conference Publication* (pp. 209-224). NASA; 1998.
- [SSH+02] Stockinger, H., Samar, A., Holtman, K., Allcock, B., Foster, I., & Tierney, B. (2002). **File and object replication in data Grids**. *Cluster Computing*, 5(3), 305-314.
- [Swi13] Switch Information Technology Services, **Description of the SLCS**.
http://www.switch.ch/Grid/slcs/about/about_long.html
- [TAA+06] Turkheimer, F. E., Aston, J. A. D., Asselin, M. C., & Hinz, R. (2006). **Multi-resolution Bayesian regression in PET dynamic studies using wavelets**. *Neuroimage*, 32(1), 111-121.
- [Tag13] The Americas Grid Policy Management Authority. **Short Lived Credential Services X.509 Public Key Certification Authorities**.
http://www.tagpma.org/authn_profiles/slcs
- [TPF+09] Torterolo, L., Porro, I., Fato, M., Melato, M., Calanducci, A., & Barbera, R. (2009, October). **Building science gateways with EnginFrame: A Life Science example**. In *Proceedings of the International Workshop on Portals for Life Sciences*, S. Gosing and J. van Hemert, Eds.
- [WFK+04] Welch, V., Foster, I., Kesselman, C., Mulmo, O., Pearlman, L., Tuecke, S., ... & Siebenlist, F. (2004, April). **X. 509 proxy certificates for dynamic delegation**. In *3rd annual PKI R&D workshop* (Vol. 14).

- [WGK+08] Wilkins-Diehr, N., Gannon, D., Klimeck, G., Oster, S., & Pamidighantam, S. (2008). **TeraGrid science gateways and their impact on science.** *Computer*, 41(11), 32-41.
- [Wil07] Wilkins-Diehr, N. (2007). Special issue: Science gateways—**Common community interfaces to Grid resources.** *Concurrency and Computation: Practice and Experience*, 19(6), 743-749.
- [WS01] Erwin, D. W., & Snelling, D. F. (2001). **UNICORE: A Grid computing environment.** In *Euro-Par 2001 Parallel Processing* (pp. 825-834). Springer Berlin Heidelberg.
- [WST10a] WSTIERIA Project Technical Note 1 (2010, April). **Federated Access to an HTTP Web Service Using Apache.**
<http://edina.ac.uk/projects/wstieria/files/TN01-facade.pdf>
- [WST10b] WSTIERIA Project Technical Note 2 (2010, June). **An Investigation of Federated Access Management for WebDav.**
<http://edina.ac.uk/projects/wstieria/files/TN02-webdav.pdf>
- [YHK95] Yeong, W., Howes, T., & Kille, S. (1995). **Lightweight directory access protocol.**