The DECIDE Science Gateway

V. Ardizzone[1,3,*], R. Barbera[1,2,3], A. Calanducci[1,3], M. Fargetta[1,3], G. La Rocca[3], S. Monforte[3], F. Pistagna[3], R. Rotondo[3], D. Scardaci[3]

[1]Consorzio COMETA, Via S. Sofia 64, 95123 Catania, Italy.

[2]Department of Physics and Astronomy of the University of Catania, Viale A. Doria 6, 95125 Catania, Italy.

[3]Italian National Institute of Nuclear Physics, Division of Catania, Via S. Sofia 64, 95123 Catania, Italy.

ABSTRACT

Motivation: The present paper reports on the architecture and the present implementation of the Science Gateway developed in the context of the DECIDE project. The motivation of the work is to enable e-Health for European citizens irrespective of their social and financial status and their place of residence, providing them with access to a high-quality early diagnostic and prognostic service for the Alzheimer Disease and other forms of dementia, based on the European research network and Grid infrastructure.

1 INTRODUCTION

The field of medical imaging has developed enormously in the past 20 years. Image databases made of thousands of medical images are now available to be used as a reference for individual diagnosis. At the same time, sophisticated and computationally intensive algorithms have been developed that can extract information from medical images invisible to the naked eye. In particular, brain diseases are ready to benefit from such applications. Highly prevalent and burdensome chronic conditions such as Alzheimer Disease (AD) and other neurodegenerative and neurodevelopmental disorders can be diagnosed early with image-based markers of structural and functional brain changes, allowing early pharmacological or rehabilitative interventions. Each year, 1.4 million Europeans will develop a form of dementia (one every 24 seconds) and it is believed that currently there are 7.3 million Europeans living with dementia and about 35.6 million people worldwide. In addition, that number is estimated to nearly double over the next 20 years to 65.7 million people in 2030. In 2008, the total cost of illness of dementia disorders in the European Union was estimated to 160 billion Euro of which 56% were costs of informal care. Because of the ageing population and increasing pressures on public finances, dementia will become one of the major challenges in the next decades for the sustainability of national health systems.

Unfortunately, neuroinformatics advancements require high computational and storage resources as well as large reference image datasets of normal persons, confining their use to advanced academic hospitals and research centres equipped with appropriate human expertise and computational facilities.

Aim of the Diagnostic Enhancement of Confidence by an International Distributed Environment (DECIDE) project [DECIDE], co-

funded by the European Union under its Seventh Framework Program, is to design, implement, and validate a dedicated e-Infrastructure relying on the Pan-European backbone GÉANT [GEANT] and the National Research and Education Networks (NRENs) and on the European Grid Infrastructure EGI.eu [EGI] and the National Grid Initiatives (NGIs) and based on the research infrastructure of neuGRID.

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

- Automated segmentation of hippocampal volume from structural magnetic resonance images to support the diagnosis of AD;
- Voxel-based statistical analysis of 18F-FDG positron emission tomography (PET) and Tc99-ECD single photon emission tomography (SPECT) to assess patterns of brain hypo-metabolism and hypo-perfusion to support the diagnosis of AD;
- Spectral-based statistical analysis of electroencephalographic studies, used for the extraction of quantitative electrophysiological markers to support the diagnosis of AD;
- Pattern recognition analysis of functional neuroimaging studies, already assessed for the extraction of class-related biomarkers in the classification of schizophrenic patients with 18FDOPA PET and extended for functional 18F-FDG-PET in neurodegenerative dementia.

DECIDE applications and tools are exposed to the end users (neurologists, physicians, and scientists in general) through a Science Gateway [Wilkins-Diehr 2007, Wilkins-Diehr 2008].

In this paper the DECIDE Science Gateway is presented from the technical and technological point of view. The paper is organized as follows. Section 2 describes the architecture of the DECIDE infrastructure and the methods and technologies used to build its application portal. Section 3 reports on the implementation done so far and the first results obtained. Conclusions are drawn in Section 4.

*Corresponding author; email: valeria.ardizzone@ct.infn.it.

2 METHODS

The DECIDE platform is built on top of three fundamental pillars: network connectivity, Grid computing resources and domain-specific scientific applications (see Figure 1). The network connectivity brings together different type of structures (clinical and research centers and academic research institutions) with a customized interconnection among all partner sites and granting high speed/large bandwidth and reliable access to the Grid infrastructure. The Grid infrastructure is used as a collaboration tool among partners as a technological glue to harmonize and unify developments and as an elastic pool of computing and storage resources where to host large volumes of data and perform their analyses. The Grid of DECIDE relies on the European GÉANT network and provides partner sites with direct links to their NRENs. DECIDE applications refer to four different diagnostic/prognostic algorithms which are based on advanced approaches to handle complex images and aim at enhancing diagnostic confidence. Neuroimaging markers will be extracted by the techniques listed in the previous section, comparing the neuroimaging data of the patients to large reference database shared by the hospitals interconnected by the e-Infrastructure. The DECIDE services will be validated in cutting-edge clinical conditions and the diagnosis of schizophrenia will also be addressed.
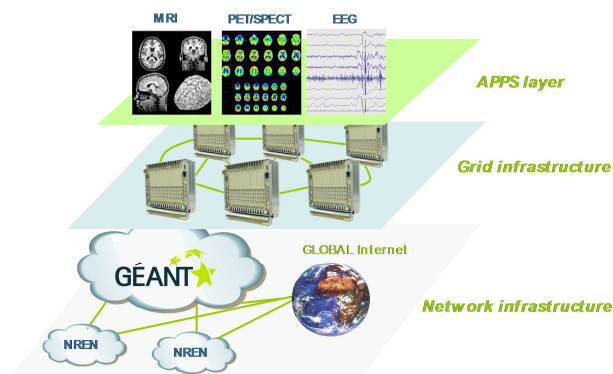

Fig. 1. Multi-layered architecture of the DECIDE platform.

DECIDE is focused on supporting neurologists and physicians involved in the assessment of neurodegenerative diseases in the diagnosis and prognosis and aims at enhancing users 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 tries to support the daily needs of neurologists while dealing with their patients, going well beyond the world of research.

The vertical approach to e-Health adopted by DECIDE ensures the requirements of the neurological community to be taken into account from the very beginning in the design of application services to ensure full usability in a real clinical environment. The use of four different medical acquisition data (Magnetic Resonance Imaging - MRI, Positron Emission Tomography - PET, Single Photon Emission Computed Tomography - SPECT, and Electroencephalography - EEG) allows combining complementary diagnostic approaches on neurodegenerative disease diagnosis, enabling synergies between different clinical domains and possibly supporting correlation studies among different clinical approaches in the field

of neurology. Four different diagnostic/prognostic algorithms are planned to be provided as services in the DECIDE Science Gateway. They are based on advanced approaches for the enhancement of diagnostic confidence and on complex images or data processing. Mainly, their goal is 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 its prognostic relevance:

- **GridSPM** [Castiglioni 2009]: 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** [Turkheimer 2006, Bose 2008]: 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** [Morra 2008]: implements an automatic algorithm for the subcortical segmentation of single-subject MRI brain images for hippocampal volume estimation, using the auto context model (ACMAdaboost) developed by LONI [LONI];
- **GridEEG** [Babiloni 2001, Babiloni 2009, Blinowska 2010]: 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.

Moreover, the project will design and implement a multimodal imaging repository, to include MRI, PET and EEG datasets and made them available for exploitation to the data analysis software at the basis of the diagnostic/prognostic service. Medical data ownership remains of the Physicians who contribute with his medical data to the medical repository, uploading data and reports with their relevant authorization rights. No free download of medical data from the DECIDE repository will be possible, but is allows external experts to use the medical data within the repository through the DECIDE diagnostic/prognostic service.

3 RESULTS

As described in the previous section, and visually explained in Figure 2, DECIDE aims to use e-Infrastructures to allow medical experts to build a production quality service, running around the clock, which allows doctors to execute algorithms on data coming from different diagnostic instruments in order to determine brain markers for the early diagnosis of AD and other forms of dementia.
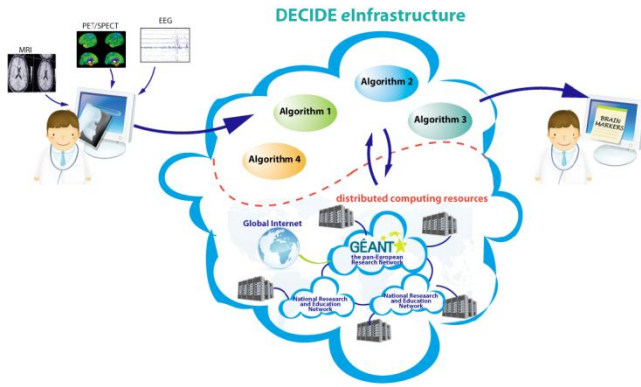
Fig. 2. Pictorial view of the DECIDE infrastructure and services.

This section describes the elements of the DECIDE infrastructure and its services and shows the results obtained so far (the project started on the 1st of September 2010). Separate sub-sections are devoted to the e-Infrastructure and to the Science Gateway.

3.1 The e-Infrastructure

As of today, the DECIDE Grid infrastructure is made of ten sites (see Figure 3). Six of them, all officially belonging to EGI, constitute the production infrastructure while four constitute the pre-production infrastructure where the algorithms are developed and tested before being fully deployed. One of the sites (FBF) is also a site of the Grid infrastructures of the neuGRID [neuGRID] project with which DECIDE will be interoperable in terms of services, data and applications.
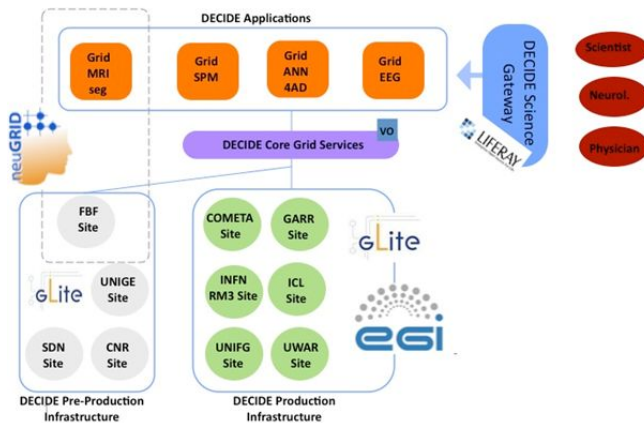

Fig. 3. Layout of the DECIDE Grid infrastructure.

On all the sites of DECIDE, the latest version of the gLite middleware [gLite] is deployed and all of its most common services are installed and running. A dedicated instance of the Virtual Organisation Membership Service (VOMS) is also available.

Besides the standard gLite middleware, two additional Grid services based on gLite are also deployed: the gLibrary framework for Grid-based digital repositories [Calanducci 2007] and the Secure

Storage System for on-the-fly data encryption/decryption [Scardaci 2007] which has proven to be robust and scalable [Scardaci 2009] and uniquely providing the requested features.

gLibrary is a robust, secure and easy-to-use system to handle wide-spread digital assets stored on a distributed Grid infrastructure. All entries in gLibrary are organized according to their type: a list of specific attributes describe each kind of asset to be managed by the system. These are the same attributes that can be queried by users. Assets are associated with the proper type in the registration/upload process. An asset catalogued as a given subtype inherits the attributes of its parent type. Of course, types are defined according to the users' needs and taking into account the assets they want to manage. The flexibility and extensibility offered by this system allow different communities to adopt gLibrary for many cataloguing purposes. Input files can be read from local disks, network shared folders, HTTP/FTP servers, etc. and replicated to one or more storage elements on which the user is authorized to write. gLibrary can also manage assets already present on Grid resources, through direct access to File Catalogues. A fine-grained authorization mechanism is used to set permissions: each asset, type and category has a set of ACLs that restricts its usage, allowing asset owners to grant access to selected groups or just a single user. Users can view in the browsing interface only those entries, types and categories for which they are granted access privileges.
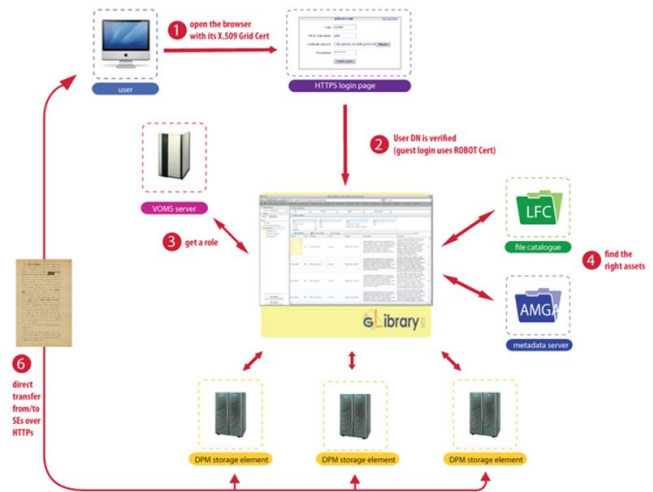

Fig. 4. Architecture of the gLibrary system.

gLibrary is built on top of the Lite middleware and uses the following services (see Figure 4), all deployed on the DECIDE infrastructure:

- The Storage Elements (SEs) that provide uniform access to data storage resources. They can be single disks, large disk arrays or tape-based Mass Storage Systems;
- The AMGA Metadata Catalogue [AMGA] that stores metadata describing the contents of Grid files, allowing users to search for entries based on their descriptions;
- The LCG File Catalog (LFC) that maps logical filenames onto the physical locations of replicas of a file stored in one or more Storage Elements;

- The Virtual Organization Membership Service (VOMS) that allows a detailed definition of users' privileges and roles according to abstract entities called "Virtual Organizations" (VOs);
- The Information Service (IS) that provides information about Grid resources and their status; in particular, the IS is used to discover the SEs available for a given VO.

Even if at the moment gLibrary is very gLite-centric, it can easuly be easily integrated with other storage technologies, such as cloud platforms, as far as they provide some kind of URL for referring to files and support common transfer protocols such as HTTP/HTTPS, FTP, GSIFTP, etc..

One competitor of gLibrary is the gCube framework (www.gcube-system.org) developed in the context of the DILIGENT and D4SCIENCE projects. gCube provides many features but at the cost of an increased complexity in the initial setup, deployment and management of repositories. gLibrary currently provides less features with respect to gCube but it does it through a very easy-to-use and intuitive interface, hiding almost completely to the users the complexity of the underlying infrastructure.

The Secure Storage System provides users with suitable and simple tools to save confidential data in storage elements owned by an external organization in a transparent and secure way, hiding the complexity of the operations necessary to ensure data privacy, integrity and availability. The core component of the Secure Storage is the keystore, a new grid element used to store and retrieve the users' keys in a seure way. The keystore has to be installed inside the data owner's trusted environment and not accessible from the external world to guarantee a good security level. The Secure Storage Service has been designed to be integrated in the gLite middleware and it is made of the following components:

- Command Line Applications: commands integrated in the gLite User Interface to encrypt and upload, decrypt and download files on the storage elements;
- An Application Program Interface: the API allows the developer to write programs able to manage confidential data using the Secure Storage service;
- The Keystore: a new grid element used to store and retrieve the users' keys in a secure way;
- The Secure Storage Framework: is a component of the service, internally used by the other components. It provides encryption/decryption functions and other utility functions. It takes care of interaction with the Grid Data Management System.

As an example, one of the Secure Storage commands is graphically explained in Figure 5.
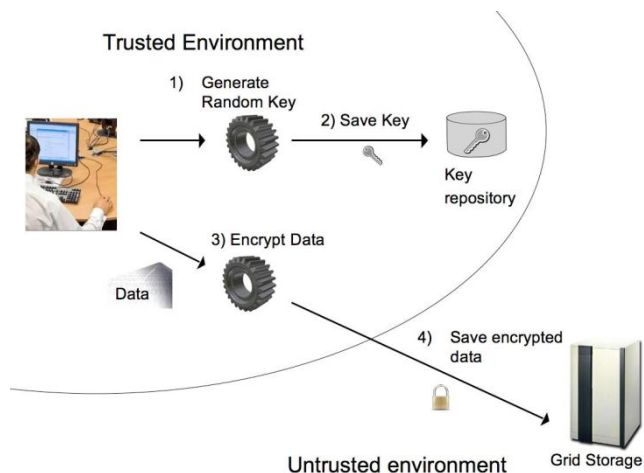

Fig. 5. Example of Secure Storage commands (lcg-scr). This command uploads and encrypts a file on a storage element doing the following actions: 1) a new random secret key is generated; 2) the key and the ACL are saved on the keystore; 3) the input file is encrypted inside user trusted environment; 4) The encrypted file is uploaded on the Grid Storage Element.

The Secure Storage service stores user files in a Storage Element in an encrypted format. An authorized user could in principle download a file from a Storage Element breaking the access policy but, in any case, he/she would not be able to decrypt it because he/she does not own the key needed to do it. Then, data access control of the Secure Storage Service is based on the policy to access the keys on the keystore. Indeed, a user needs to get the proper decryption key from the keystore to access data in a clear format.

The Secure Storage Service authorization model has been designed to be integrated in the gLite middleware using the standard credentials (proxy certificates with VOMS extensions) used in this environment. In this way, users can exploit Secure Storage using their gLite credentials without the need to install new security software. The keystore implements an authentication procedure based on the information stored in the user's proxy (user Distinguished Name and VOMS attributes). It provides or denies the key needed to decrypt the data using an Access Control List (ACL) mechanism. An ACL is associated to each decryption key and it can be made of one or more distinguished names (DNs) and/or one or more VOMS attributes. It extracts the DN and VOMS attributes from the X.509 proxy certificate and checks if the user is authorized. The keystore provides users with the decryption key only if their DNs or VOMS attributes contained in their proxy match with an entry in the ACL of the key.

3.2 The Science Gateway

This section describes the architecture and present status of the DECIDE Science Gateway. As shown in Figure 3, the Science Gateway is built within the Liferay framework and container [Liferay] and it is fully compliant with the JSR 268 ("portlet 2.0") standard. Separate sub-sections are devoted to the various functional aspects of the portal.

3.2.1 Authentication and Authorization

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 (Grid non-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 still 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 fulfil the stringent requirements and implements the expected roles and corresponding privileges.

Moreover, in order not to confuse inexperienced users with different sets of credentials, another design requirement was to have in place a Single Sign On (SSO) mechanism across all services a given user is entitled (i.e., has the right) to use.

The above requirements have been fulfilled by the adoption of the Shibboleth System [Shibboleth] for authentication and the Security Assertion Markup Language (SAML) to implement the SSO. Shibboleth allows institutions wishing to include the DECIDE Science Gateway as one of the resources of their users to simply and easily create an Identity Provider (IdP). When a user tries to use one of the DECIDE applications available on the Science Gateway, he/she is re-directed to the IdP of his/her own institute and the IdP is responsible for the identification of 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 which the user is automatically logged in.

Currently, the portal is part of GrIDP federation, a new federation operated by Consorzio COMETA to manage several web portals. Nevertheless, a formal request to join the IDEM federation [IDEM], one of the biggest Shibboleth federations available. provided by GARR, and including many Italian universities and research centres, has also been submitted.

Once a user is authenticated, the authorisation system verifies his/her credentials and the Scientific Board of DECIDE grant authorisations. A centralised LDAP server provides the authorisations by associating users with roles so a user can perform on the Science Gateway all the activities designed for the roles he/she is associated with.

Once the user is authenticated and authorised to run one the DECIDE applications, the last step to be done is the creation of a proxy certificate to secure Grid transactions. Usually, this requires the user to have a personal X.509 digital certificate and be registered in the VOMS of a given Virtual Organisation. Furthermore, he/she also has to have his/her certificate loaded in the web browser which is very often a solution prone to security breaches. The adoption of personal certificates to access e-Infrastructures has demonstrated to be difficult by non-expert users and represents a limiting factor to the rapid spreading 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 as smooth as possible, has recently been achieved with the introduction of robot certificates, also referred as portal certificates. The advantages introduced by this new kind of digital certificates are manifold and they have currently been adopted by several Certification Authorities such as those of UK, The Netherlands, and Italy. Robot certificates are nowadays successfully used, for instance, to automate Grid service monitoring, distributed data collection systems, and identify a responsible for unattended services one wants to share with all the members of a specific VO. From a security point of view, 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.

In order to let physicians involved in the DECIDE project to access the computing and storage Grid resources through the Science Gateway, a new Grid authentication mechanism based on the use of robot certificates available on smart cards has been designed. The solution implemented (see Figure 6) extends the native Java™ Cryptographic Token Interface Standard (PKCS#11) [PKCS#11] with the Java CoG Kit [von Laszewski 2001] and the Bouncy Castle [Bouncy Castle] APIs in order to implement a "lightweight" crypto-utility which may be used by generic Grid users, client applications, Grid portals and/or Science Gateways to access robot certificates stored on smartcards and generate a proxy with VOMS extensions.
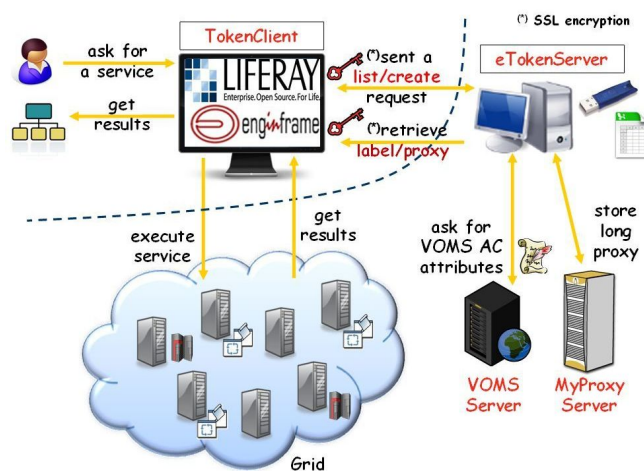

Fig. 6. Overview of how the new "lightweight" crypto-library works.

The core of the new library is represented by the eTokenServer Java class, a multithreaded server which accepts all the requests coming from a list of authorized clients and manages a list of robot certificates kept in the USB token. The client requests are satisfied by the TokenClient Java class. With this class, users, client applications, Grid portals and/or Science Gateways can send requests to the eTokenServer for browsing the available X.509 certificates or generate Grid proxies with VOMS extensions. To improve the security between clients and server, the SSL protocol is used to secure the communications.

Using this library it is possible to grant different VO attributes (roles and privileges) to the user depending on the application/task he/she wants to execute. The association of this grant is handled by the Science Gateway which takes care of providing the users with a valid temporary proxy.

The main difference with Grid portals available in other projects is the use of two different security systems linked together by the portal, providing users with an easy access to resources without the need of personal certificates. From a security point of view, the

authentication method is delegated to the institutions that can implement very restricted approach. It is also possible to have even better authentication methods than PKI certificates, e.g. mixing different approaches like password, biometrical, IP and so on. Additionally, the communication between the IdPs and the portal is encrypted so the authentication step provides a security level at least comparable with other approaches.

On the other hand, the LDAP-based authorisation allows users to use the services provided by the portal. Actually, users cannot access the resources but they have to demand to specific components the communication with the services. Since users cannot access without Shibboleth-based verification and the available services do not provide direct access to resources, it is almost impossible for users to perform malicious operations through the portal.

However, in order to avoid any abuse, a pro-active logging system registers all users' activities and matches these with the jobs registered in the gLite Logging and Bookkeeping (LB) service. This information allows identifying all the operations ensuring the non-repudiability of Grid transactions which is one of the fundamental requirements of the Grid Security Infrastructure (GSI).

Finally, the global security mechanism provides a safe environment, at least comparable to a full PKI, where medical data can be managed without security or confidentiality problems.

3.2.2   Interface to Grid services

Once authenticated to the Science Gateway, and authorized to run one of the DECIDE algorithms, users can choose one of the applications and start the procedure to submit an analysis job. The typical scenario that has been agreed with the physicians working in the project is the following:

- The user fills a web form on the Science Gateway defining the input parameters of the application;
- Input files to be analyzed by the selected algorithm are transferred to the Science Gateway;
- A job, described using the Job Description Language of gLite, is automatically created and submitted to the DECIDE Grid infrastructure together with the input files;
- The user is notified when the job is submitted and from then on he/she can monitor its status through a dedicated portlet of the Science Gateway;
- When the job finishes, the user receives an email from the Science Gateway containing the output of the job.

The back-end engine that implements the above described scenario and interacts with the gLite Grid services behind the Science Gateway front-end has been written in pure Java using the jLite API [jLite] called through the functions of the jSAGA library [jSAGA]. jLite is a Java library providing simple API for accessing gLite-based Grid infrastructures. It is intended for Java developers who would like to avoid dealing with the complexities of the gLite middleware and want to reduce time and effort needed to build cross-platform Grid applications. jSAGA is a Java implementation of SAGA (Simple API for Grid Applications) [Goodale 2011] standard specification defined by the Open Grid Forum [OGF]. jSAGA:

- Enables uniform data management and execution management across existing grid infrastructures;
- Makes extensions easy: adaptor interfaces are designed to minimize coding effort for integrating support of different middleware (besides gLite, the Globus Toolkit [Globus] and UNICORE [UNICORE] are also supported);

- Ensures operating system independency: most of the provided adaptors are written in pure Java and are tested both on MS Windows and Linux operating systems.

As shown in Figure 7, middleware interfaces are exposed to end users through standard portlets embedded in the Liferay container. Grid transactions are secured by proxy certificates created by the robot server described in the previous sub-section while data management services are used through the Representational State Transfer (REST) functions of the gLibrary framework described in Section 3.1.
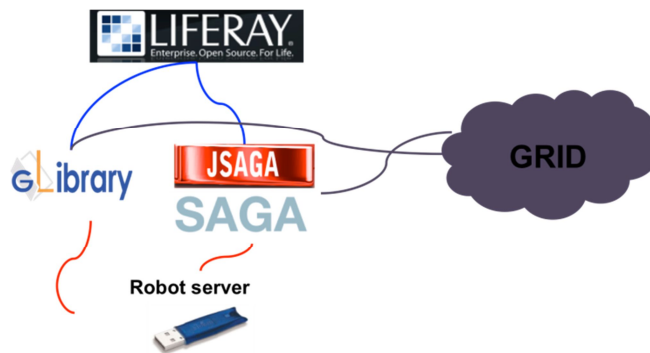

Fig. 7. Schematic view of the main components of the DECIDE Science Gateway.

3.2.3   User interface

As already mentioned above, the graphic front-end of the DECIDE Science Gateway has been developed using the Liferay portal framework and portlet container. 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. Other 200+ portlets are available in the repository of the Liferay community.

As an example, Figure 8 shows the input page of the GridSPM application available on the DECIDE Science Gateway.

Fig. 8. Input page of the GridSPM application available on the DECIDE Science Gateway.

To submit a job, users just have to select the patient gender, insert the patient age, select the input images and... click a button.

Figure 9 shows a portlet that reminds the input parameters and shows the status of the submitted jobs.

When a job ends, the user is notified by email and the output is sent to him/her as an attachment. Figure 10 shows the notification email and one example of job output.

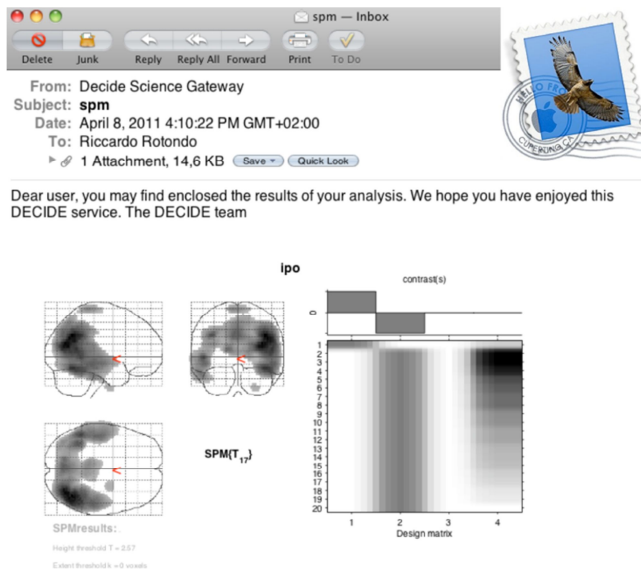**Fig. 9.** Input data and status monitoring portlets of the GridSPM application.

Fig. 10. End-job notification email and example of output of the GridSPM application.

4   SUMMARY AND CONCLUSIONS

The main goal of the DECIDE project is to 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 GEANT 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 address new challenges in their domain.

The service that will be realized by the DECIDE project will be exposed to end users as a Science Gateway based on the Liferay portlet container and the gLite middleware and makes use of sophisticated authentication and authorization mechanism able to ease the access and use still implementing a fine grained control on roles and corresponding privileges. The DECIDE Science Gateway will allow the creation and management of large distributed repositories of medical images with the possibility to encrypt the stored data.

The sustainability of DECIDE, at level of infrastructure, is ensured by the fact that all sites forming the production infrastructure belong to organisations which are members of the National Grid Initiatives established in their countries. At user lever, different initiatives have been envisaged and already planned to reach long term sustainability. Examples are the training courses, for the accurate use of the DECIDE services, which will be provided during the lifetime of the project.

ACKNOWLEDGEMENTS

REFERENCES

DECIDE. All information about DECIDE, including deliverables and activity reports, can be found on the project website: www.eu-decide.eu.

GEANT. All information about GEANT can be found on the website www.geant.net.

EGI. All information about the European Grid Infrastructure can be found on the website www.egi.eu.

LONI. All information about the Laboratory of Neuro Imaging of the University of California at Los Angeles can be found on the website www.loni.ucla.edu.

Wilkins-Diehr N. (2007), Special Issue: Science Gateways - Common Community Interfaces to Grid Resources, *Concurrency and Computation: Practice and Experience, 19(6), 743-749.*

Wilkins-Diehr N., Gannon D., Klimeck G., Oster S., Pamidighantam S. (2008), TeraGrid Science Gateways and Their Impact on Science, *IEEE Computer 41(11), 32-41.*

Castiglioni I. et al. (2009). A Grid-based SPM service (GriSPM) for SPECT and PET neurological studies, *Eur J Nucl Med Mol Imaging 36*, 1193–1195.

Turkheimer F.E., Aston J. A. D., Asselin M. C. and Hinz R. (2006) Multi-resolution Bayesian regression in PET dynamic studies using wavelets, *NeuroImage 32*, 111-121.

Bose S. K., Turkheimer F. E., Howes O. D., Mehta, M. A. , Cunliffe R., Stokes P. R., Grasby P. M. (2008). The application of an artificial neural network to classification of schizophrenic patients and healthy controls using [18F] Fluorodopa PET imaging, *Schizophrenia Research 106(2-3)*, 148-55.

Morra J. H., Tu Z., Apostolova L. G., Green A. E., Avedissian C., Madsen S. K., Parikshak N., Hua X., Toga A. W. , Jack C. R. Jr., Weiner M. W., Thompson P. M., Alzheimer's Disease Neuroimaging Initiative (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, Erratum in *Neuroimage 44(4)*, 1439.

Babiloni F., Carducci F., Cincotti F., Del Gratta C., Pizzella V., Romani G. L, Rossini P. M., Tecchio F., Babiloni C. (2001). Linear Inverse Source Estimate of Combined EEG and MEG Data Related to Voluntary Movements, *Human Brain Mapping* **14**, 197–209.

Babiloni C. et al. (2009). Directionality of EEG synchronization in Alzheimer's disease subjects, *Neurobiology of Aging* **30**, 93-102.

Blinowska K., Kus R., Kaminski M., Janiszewska J. (2010). Transmission of brain activity during cognitive task, *Brain Topography* **23**, 205-213.

neuGRID. All information about neuGRID, including deliverables and activity reports, can be found on the project website: www.neugrid.eu.

gLite. All information about the gLite middleware, including specification documents and user guides, can be found on the website glite.cern.ch.

A. Calanducci et al. (2007), **"*A Digital Library Management System for the Grid*"**, Fourth International Workshop on Emerging Technologies for Next-generation GRID (ETNGRID 2007) at 16th IEEE International Workshops on Enabling Technologies: Infrastructures for Collaborative Enterprises (WETICE-2007), GET/INT Paris, France, June 18-20, 2007

Andronico G. et al. (2011), *e-Infrastructures for Cultural Heritage Applications*, Handbook of Research Technologies and Cultural Heritage: Applications and Environments", G. Styliaras, D. Koukopoulos, F. Lazarinis (Eds.), IGI Global 2011, ISBN 978-1-60960-044-0 (hardcover) and 978-1-60960-045-7 (ebook), p. 341-369.

AMGA. All information about the AMGA Metadata Catalogue can be found on the website amga.web.cern.ch/amga.

Scardaci D., Scuderi G. (2007), *A Secure Storage Service for the gLite Middleware*, Proceedings of the Third International Symposium on Information Assurance and Security, p. 261-266.

Scardaci D. (2009). *Using Secure Storage Service inside the EELA-2 Infrastructures,* Proceedings of the Second EELA-2 Conference, p. 167-173.

Liferay. All information about the Liferay portal framework can be found on the website www.liferay.com.

Shibboleth. All information about the Shibboleth System can be found on the website shibboleth.internet2.edu.

IDEM. All information about IDEM Federation can be found on the website www.idem.garr.it/.

PKCS#11. All information about the Cryptographic Token Interface Standard, including the specification documents, can be found on the website www.rsa.com/rsalabs/node.asp?id=2133.

von Laszewski G., Foster I., Gawor J., and Lane P. (2001), A Java Commodity Grid Kit, *Concurrency and Computation: Practice and Experience*, **13(89)**, 643-662. http://www.mcs.anl.gov/~gregor/papers/vonLaszewski--cog-cpe-final.pdf.

Bouncy Castle. All information about Bouncy Castle APIs can be found on the website www.bouncycastle.org.

jLite. All information about jLite APIs, including source code and user manual, can be found on the website code.google.com/p/jlite.

jSAGA. All information about jSAGA, including source code and user manual, can be found on the website grid.in2p3.fr/jsaga.

Goodale T. et al. (2011), A Simple API for Grid Applications (SAGA), www.ogf.org/documents/GFD.90.pdf.

OGF. All information about the Open Grid Forum and its activities and standards can be found on the website www.ogf.org.

Globus. All information about the Globus Toolkit, including source code and user manuals, can be found on the website www.globus.org.

UNICORE. All information about the UNICORE middleware, including source code and user manuals, can be found on the website www.unicore.eu.