2012 Sixth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing

A Service Oriented Analysis Environment for Neuroimaging Studies

Ashiq Anjum, Nik Bessis
School of Computing and Mathematics
University of Derby, Kedelston Road Campus, Derby, UK, DE22 1GB
{a.anjum, n.bessis}@derby.ac.uk

Richard McClatchey, Kamran Munir, Irfan Habib, Andrew Branson, Jetendr Shamdasani
Centre for Complex Cooperative Systems
University of the West of England, UWE
Bristol BS16 1QY UK
{Richard.McClatchey, Kamran.Munir, Irfan.Habib, Andrew.Branson, Jetendr.Shamdasani}@cern.ch

***Abstract*— neuGRID is an EC-funded project driven by the needs of the Alzheimer's disease research community that aims to facilitate the collection and archiving of large amounts of imaging data coupled with a set of generalized services and algorithms. By taking Alzheimer's disease as an exemplar, the neuGRID project has developed a set of generic analysis services and a Grid infrastructure that enables the European neuroscience community to carry out research required for the study of degenerative brain diseases. Using the services and the infrastructure, neuroscientists should be able to more easily identify neurodegenerative disease markers through the analysis of magnetic resonance and other brain imaging. The availability of such image-based disease markers will allow earlier diagnosis and foster the development of new drugs. This paper reports our work on the service oriented analysis environment that has been developed from a study of user requirements and that enables the neuroscience community to conduct and trace analyses for the study of Alzheimer's and other neurodegenerative diseases. We present the salient features, architecture and implementation details of the services that will form an analysis environment. We also describe the functionality and benefits that these services will offer to the medical community in general and neuroimaging analysis in particular.**

Keywords-SOA; neuroimaging Analysis; Grid Computing; Workflows

I. INTRODUCTION

Alzheimer's disease is a progressive, degenerative and irreversible brain disorder that causes intellectual impairment, disorientation and eventual death. It is the most common cause of dementia, accounting for around two thirds of cases in the elderly. It is estimated that 2-5% of people over 65 years of age and up to 20% of those over 85 years of age suffer from the disease [1]. The study of Alzheimer's disease, its causes, its symptoms and especially its early diagnosis is now a major driver in the provision of healthcare for the elderly. Early diagnosis is beneficial for several reasons. Having an early diagnosis and starting treatment in the early stages of the disease can help preserve function for months to years. Having an early diagnosis also helps families plan for the future, make living arrangements, take care of financial and legal matters, and develop support networks.

Alzheimer's clinical researchers are currently seeking the assistance of large-scale information technology resources to enable them to study masses of neuroimaging data being accumulated across the older patient community so that early onset indicators such as cortical thinning can be studied [2]. Rapid advances in neuroimaging technologies such as PET, SPECT, MR spectroscopy, DTI and fMRI have offered a new vision into the pathophysiology of AD [3] and, consequently, increasingly new powerful data analysis methods have been developed [4]. Since the beginning of the new century the development of innovative techniques for ROI-based volumetry, automated voxel based morphometry, cortical thickness measurement, basal forebrain volumetry, and multivariate statistics have emerged [5, 6]. Methods which were felt to be the most feasible and accurate have started to be used in clinical settings. The availability of large image data repositories to the neuroimaging community has necessitated the development of distributed data and processing infrastructures to access data and online image analysis tools and to assess longitudinal brain changes [7, 8, 9].

Many efforts have been directed at creating brain image repositories including the recent US Alzheimer Disease Neuroimaging Initiative (ADNI). Multi-site distributed computing infrastructures have been launched with the goal of fostering shared resources and intensive data analysis to advance knowledge about neurodegenerative diseases. Numerous efforts, such as NeuroLOG [10] and Neurogrid [11], have been conducted which focus on providing grid infrastructures that support neuroimaging applications [12]. At present, however, these applications tend to be either focused on specific pathologies or are directed at supporting a subset of neuroimaging applications. Moreover, these solutions are tightly bound to specific platforms, which may limit their wider adoption across neuroscience. neuGRID is an effort which targets the limitations of existing neuroimaging based Grid infrastructures. It aims to provide an infrastructure and a set

978-0-7695-4684-1/12 $26.00 © 2012 IEEE
DOI 10.1109/IMIS.2012.140

of complementary analysis services that are designed to support and enhance research, which is necessary for the study of neurodegenerative diseases. neuGRID is an EC-funded effort, arising from the needs of the Alzheimer's disease (AD) imaging community, which will allow the collection and archiving of large amounts of imaging data paired with services, Grid-based algorithms and computational resources. The major benefit will be faster discovery of new disease markers that will be valuable for earlier diagnosis and development of innovative drugs.

In this paper we present work on the set of generalized services in the neuGRID project that has been specified in consultation with its clinician user communities and developed to facilitate neuroimaging analysis, such as Alzheimer's studies. A set of services has been developed using the service oriented architecture paradigm that can in principle be ported and customised for other related medical domains. The services can thereby be reusable both across Grid-based neurological data and later for wider medical analyses. The approach used to transform the medical users' requirements into discrete Grid services is also outlined in this paper. The services include workflow (or analysis pipelines), provenance, querying, glueing and abstraction services. The services can 'glue' a wide range of user applications to a range of available Grid platforms thereby creating a foundation of cross-community and cross-platform services for neuroimaging analysis and promoting interoperability between diverse projects in this domain. neuGRID aims to develop an infrastructure that will adhere to both Web Service and relevant standards during the development of portable high-level Grid services. Such facilities will allow clinicians to query, analyse, track, annotate and preserve data that is resident on the Grid.

II. Neuroimaging Analysis in neuGRID and Alzheimer's Disease

Brain imaging is currently one of the fastest expanding areas of biomedical interest in the search for biomarkers of brain diseases. Whereas the accessibility of other organs has facilitated the development of biomarkers (e.g.the diagnosis and monitoring of a cancer through specific proteins present in the bloodstream), the brain is a delicate organ and Nature has thought fit to protect it tightly from outside influences – even those coming from the body itself.–Not only is the brain protected from chemical damages by the blood-brain barrier, but also from physical damage by the thick bone of the skull – the hardest and most shock resistant bone in the human body.

For these and other reasons, the development of brain disease biomarkers has lagged behind that of other body organ diseases. Brain imaging techniques may help develop markers for neurological diseases. This has already been the case with multiple sclerosis, where MR allows the recognition of typical brain lesions. This has made possible the assessment of the effectiveness of the only drug that so far has been found to delay the progression of this dreadful disease (interferon 1-beta). Other brain diseases are unfortunately much more prevalent than multiple sclerosis and for which presently we have no satisfactory treatment, including Alzheimer's disease. Once believed to be an inescapable companion of older age, in the past 10 years we have learned that it is caused by "poisoning" of the brain by a small protein (Abeta). The increased pathophysiological knowledge has led to the development of antiamyloid drugs that are now in the pre-marketing phase of human testing.

It needs to be stressed that some of the presently available algorithms can take as long as 8 hours per brain to run on a state-of-the-art workstation. The *modus operandi* today is that of scientists physically migrating image data to remote imaging centres (normally on DVDs) where they can find expertise and computational facilities for analysing small personal datasets (a few hundreds of images at most). This is the case not only for neuroimaging research partners in the neuGRID project but also for countless other scientists. Typically, a research fellow can spend months at an image analysis centre – often thousands of miles from home – where he/she learns the use of the algorithms on personal image data, then returns to the original research group, where he/she can install all or part of the procedure and run jobs either in house or remotely on the image analysis centre servers. This scenario is becoming unsustainable and it needs to change radically in the near future, when thousands of images will be available and many more researchers will want to perform image analysis.

In the US the largest project to date in the field of neurodegenerative disorders is being carried out (the North American ADNI) where persons with early onset Alzheimer's disease are being assessed every six months for four years with a number of imaging techniques in order to collect over 5000 brain images for study. In Europe, the feasibility of the adoption of the ADNI platform has been investigated in the FP6 ENIR (Foresight Study for the Development of a European NeuroImage Repository), and the data collection study co-funded by the US Alzheimer's Association (pilot European ADNI), and the FP6 Innomed-AddNeuroMed. AddNeuroMed will study 900 persons at baseline, three, and 12 months and allow to collect 2,500 brain imaging studies. Notably, these 2,500 will be fully compatible with the ADNI 7,500 images, thus reaching a dataset of 10,000 brain imaging studies in total. Moreover, the ADNI data is public, in that access is granted to any scientist wishing to exploit it.

The combination of larger datasets and a larger scientific community will make the current paradigm unsustainable in the near future and will make research environments with archiving as well as computational facilities absolutely necessary. Computational Grid-based solutions might be the answer to the archiving and computational needs of image neuroscientists. Indeed, Grid-based solutions to share data and carry out advanced computational analyses are spreading quickly in the biomedical field. These considerations led us to develop the set of analysis services and a Grid infrastructure in the neuGRID project.

III. A Service Oriented Analysis Environment in neuGRID

In order to facilitate analysis, querying and collaboration that can address the users requirements, a service oriented analysis environment has been proposed in neuGRID in which high-level distributed services such as querying, workflow management, provenance, and anoymization services[13, 14] coordinate and interact to support user analyses. Such services will help the users in sharing data and knowledge and should enrich medical decision support systems. A widely used and standard form of a generalized distributed environment which can support a wide range of applications and platforms is a so-called service oriented architecture (SOA). The service layer in neuGRID is implemented using the SOA paradigm in order to have a flexible and reusable medical services layer which can be customized for various applications. The services in the neuGRID analysis environment, coordinate to facilitate the neuroimaging analysis process, using SOA principles.

The first action in the neuroimaging analysis cycle, shown in figure 1, is to register images in the neuGRID store that have been collected from the hospital data acquisition systems or have been imported from other research projects. This analysis has been derived from a rigorous study of the requirements of clinical researchers. (The border in figure 1 denotes the limit of the neuGRID project infrastructure.) As an example, consider a new clinical site that may wish to make use of the neuGRID infrastructure to share data within a wider research community. Existing data would be put through a process which enforces quality control, formatting and ethical compliance. The data is then integrated with the neuGRID standard data model, which enables other researchers to access it and carry out their research. As new data sets are acquired they go through a local quality control step before passing through the same system-wide quality control, formatting, ethical compliance and data model integration processes that the pre-existing data goes through.

Once the data has been imported into the neuGRID system and users are able to access and query these data, they may want to carry out studies and data analyses to find results of interest. Workflow development is a methodology that can be used to represent user preferences for an automatic analysis of data; this is an important process in the analysis chain. Users may create workflows and then quickly execute the workflows on distributed resources provided by the Grid. Workflow development and execution is an important stage in the analysis life cycle in the neuGRID project. For example, a researcher may wish to run a comparative analysis using a study set of 3000 MRI scans stored in geographically distributed medical centres. It is important that the results are generated in a timely fashion since the researcher may have a number of different studies to carry out that week. The user would interact with the system to choose a study set of 3000 images, would select the pipeline or workflow through which the analysis will take place and would start the analytical process.

A service oriented analysis environment can enable the construction of a catalogue of reusable services that can be customized and reused across domains and applications with minimal change. The users can use part of, or whole services to conveniently build their applications by exploiting well defined sets of interfaces that come with each service. These services are not tied to a particular middleware and can in principle be seamlessly integrated, deployed and used in different applications as if they were specifically developed for the application. This is an essential element in making the services extensible and available for interoperation between systems. A service oriented environment also facilitates service maintenance. By focusing on the service layer as the location for core business logic, maintainability improves since developers can more easily locate and correct defects in the services. In addition services have published interfaces that can be easily tested and validated independently from any application that uses the service. The architecture itself is flexible, new services can be developed in future and embedded into the architecture with little or no modification to any other services, as they are independent, modular and loosely coupled. Finally because of standardized interfaces, the internals of the service can be developed on any platform, providing a platform- and middleware-agnostic solution. The neuGRID services have been built according to the architecture and philosophy presented in section 5 in order to enhance and promote their re-usability in other related applications. In neuGRID the services have been implemented and on top of the gLite middleware [15].

IV. Architecture and Philosophy

Most (Grid) services available to biomedical researchers have been designed for a particular community of medical users and these services have also been developed for and deployed on a particular (Grid) middleware. Services built by one community often cannot be shared or re-used in other medical domains due to particular architecture, interface or platform limitations. In order to enhance the reuseability of the neuGRID services, one of the major design considerations was to develop them in a manner that keeps them independent from the underlying Grid middleware. The neuGRID services have been designed to be middleware agnostic and to hide the heterogeneity of underlying distributed resources through a common abstraction layer.

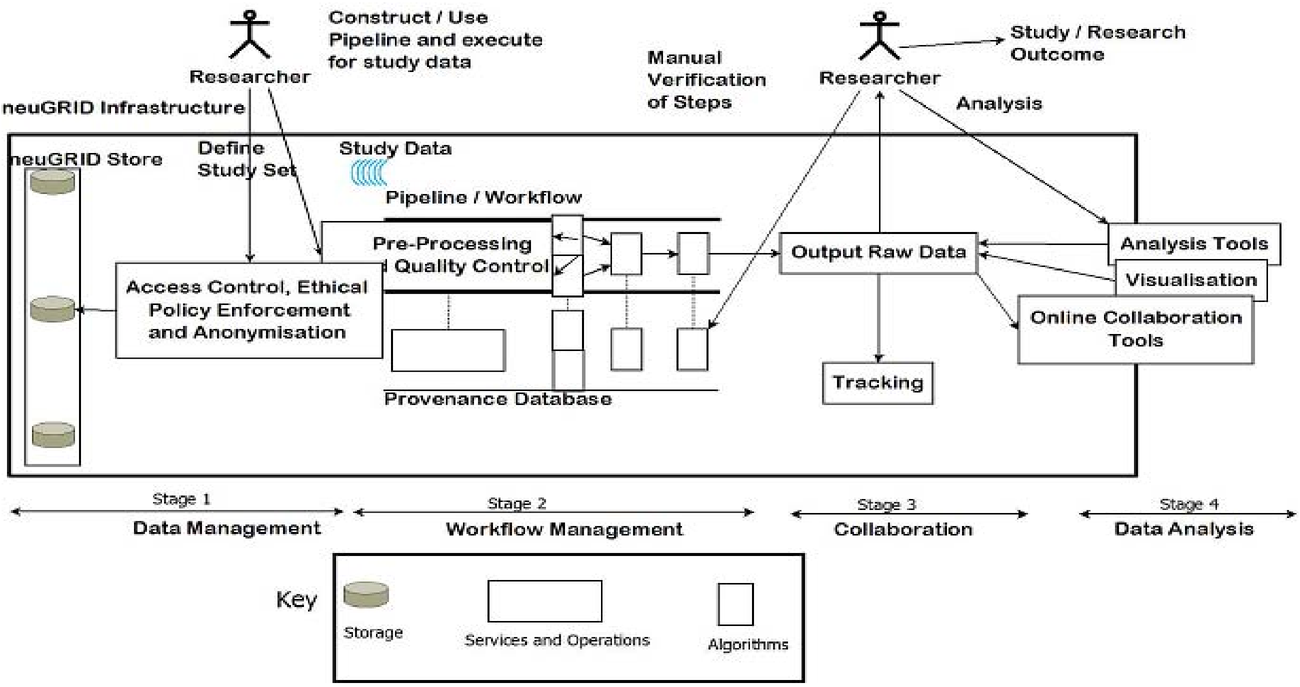

Figure 1 : An end to end example of the neuGRID analysis environment

As figure 2 illustrates, in neuGRID the distributed medical services sit between the Grid services and the user facing tools. The Grid is like a "black box" for the services and a so-called 'Glueing' service (as discussed in later sections) has been implemented to provide the required abstraction from the Grid resources. The generalized services will access Grid services through the abstraction layer in a transparent fashion in order to avoid lock-in to a particular Grid middleware. This loose coupling philosophy between the Grid and generic services is a fundamental design principle in the project.

To justify the design philosophy, consider a service previously deployed using the Grid middleware gLite [15] and a user wanting to use a different Grid. In the current scenario, this transition from gLite to another middleware is not straightforward and requires quite some changes in the code, recompilation and redeployment on the new middleware. This can divert the user's attention from her application to the infrastructure and as a result, users can become Grid-phobic or at least Grid-resistant. If a mechanism is developed whereby users are not concerned about the Grid fabric and functionality (with Grid details remaining abstracted), this should help not only to make the Grid use more attractive, but also to enable the users to concentrate on their research analyses. Applications and services developed for one platform can then be deployed on any other Grid middleware without significant user effort. This will also enable neuGRID to make use of existing resources that are running a range of middleware. This is one of the core design objectives of the neuGRID project where generalized services have been developed that can run on any middleware (see figure 3).

In addition to the abstraction from specific technologies, the main characteristics of the services are the loose coupling between them and service extensibility. This loose coupling implies that services do not invoke each other directly but rather interactions are managed by a so-called 'orchestration entity'. Technological abstraction is obtained by using service contracts that are platform-independent. Extensibility is achieved through service discovery and composition at execution time. Whilst reusability is one of the most important criteria for determining which services to build, there are other criteria that have been considered to meet the constantly changing requirements of research environments. In particular, the evolvability of services is a criterion that has been accommodated. The services have been designed in a format which enables reuse as well as composition of services in workflows. A reuseable service is one which can be reused in a number of ways, including a range of different applications, but the service consumers are application users. A composable service on the other hand, is a service which can be used in workflows and the service consumers are services themselves. Hence in order to build composable services, interfaces have been designed in a format which make them compatible with other services. Not all services are suitable for composition into workflows. Therefore, service reusability and composition had to be addressed at the same time in the service design.

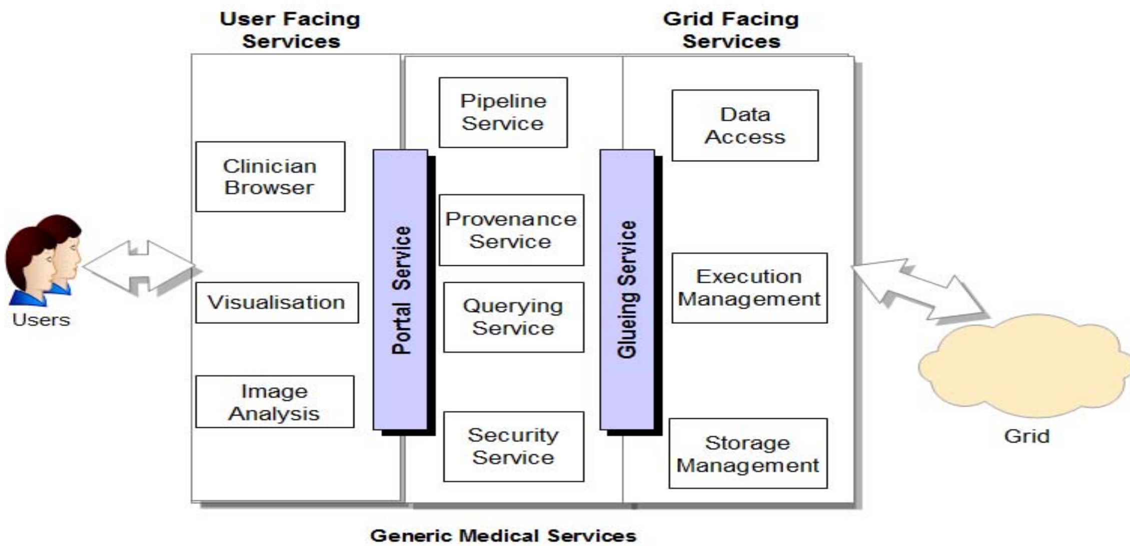

Figure 2: The neuGRID layered services architecture

V. NEUGRID SERVICES AND THEIR FUNCTIONALITY

This section describes in greater detail the services that make up the neuGRID iinfrastructure. User requirements have been distilled into a set of services that help neuGRID provide an enabling analysis environment to the neuroimaging community. The services cover each requirement depicted in the end to end requirements diagram shown in figure 1. The services are self contained and loosely coupled entities that exist independently and, in cooperation with other services, can support the user analysis process. As stated earlier, the services are divided into three groups: i) User facing Services ii) Analysis Services and iii) Grid facing Services. A description of each group of services is discussed in the following paragraphs.

The user facing services include those that are accessed by a user for her day-to-day activities. They provide the interfaces that are necessary to enable the user to leverage the underlying neuGRID services to support her analyses. The first of these services is the Portal Service. The user will interact with the system through the Portal Service which encapsulates and abstracts the complexity of the underlying neuGRID services.. This service is supported by a single sign-on enabled authentication service that enables the user to access the underlying neuGRID and Grid services without repeated authentication. This in turn is managed by a Security Service that is embedded in the Portal Service. This service is responsible for all the authentication, authorisation, access controls and policy enforcement issues. The Security Service, in association with an anonymisation service, is also responsible for the anonymisation, privacy protection and data format conversion of the datasets (for example from MINC [16] to DICOM [17]).

The next set of services shown in the services architecture (see figure 4) are designed to provide general purpose analysis facilities. This set of services, the main focus of this paper, focusses on providing functionality for managing and executing workflows, querying and managing provenance information as well as facilitating information querying and retrieval from both image datasets and other clinical data. They have been implemented in such a way that a variety of applications and Grid middleware can be supported. The Glueing Service provides the required abstraction from the underlying Grid middleware. The first of the generic services is a workflow specification and transformation service called the Pipeline Service. Through this service users can specify their pipelines (or workflows, connected collections of algorithms). This service is designed to support all major neuroimaging workflow authoring environments (such as LONI Pipeline [18], Kepler [19] etc) and with the help of the Glueing Service can enable their enactment on a number of underlying infrastructures. The Provenance Service, as depicted in the services architecture diagram, can capture steps in an analysis workflow and store the workflow and any associated meta-data generated during its enactment. A workflow management system supports the specification, execution, re-run, and monitoring of scientific processes. In such a complex workflow specification and execution it is possible for user errors to be introduced which cumulatively could have a large impact on the validity of the results that are produced. Researchers therefore require a means of tracking the execution of given workflows so they can verify that important results are accurate and reproducible. The Provenance Service is primarily intended to capture workflow specifications, final and intermediate datasets produced during an analysis as well as auxiliary information pertaining to the environment where the data was produced.

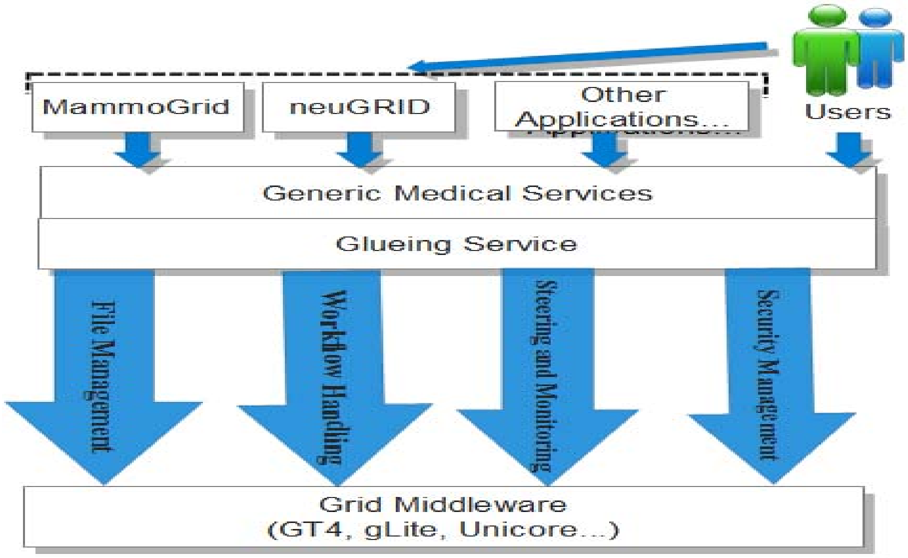

Figure 3: Middleware agnostic services

The Provenance Service allows users to query analysis information, to regenerate analysis workflows, to detect errors and unusual behaviour in past analyses and to validate analyses. The service supports the continuous fine-tuning and refinement of pipelines in the neuGRID project. After the execution of a workflow all the information that was initially provided and that which was generated during an analysis is stored in the Provenance store. This store can be queried by the user to verify results or improve and fine-tune pipelines and acts as a rich knowledge base for users.

The Querying Service provides methods to enable the efficient querying of heterogeneous data in neuGRID. This includes data formats that range from images, flat files and relational databases to XML structures. The primary data sources in neuGRID comprise clinical data sets that include images and associated metadata. These data sets are managed by the LORIS system [20]. The Glueing Service is the gateway through which all of the analysis services communicate with the underlying Grid Infrastructure. The neuGRID Glueing Service supports an architecture, which hides the heterogeneity of distributed resources and enables the services to access resources without tying them to a specific Grid middleware. It provides the necessary interfaces that enable enactment of workflows in various distributed infrastructures, querying of Grid databases as well as access to provenance information. Hence the Glueing service hides the encapsulation of Grid middleware complexities from the neuGRID services. Using the Glueing service neuGRID services can be deployed on various Grid middleware that includes gLite, Globus [21], Unicore [22] or any other SAGA [23] supported Grid middleware, thus promoting interoperability. neuGRID is based on an SOA paradigm with each component in the infrastructure being a web service. SOAP is used for communication between disparate services; the SOAP stack adopted for the project is the Apache Axis 2 Framework [24]. The Apache Tomcat web services container is used to host the web services; it is a mature, open source and extensible web services container environment. The basic single sign-on infrastructure used in the neuGRID project is the Central Authentication Server (CAS) which is a widely used Open Source SSO implementation in Java. To extend the single sign-on facility to the Grid environment, CAS has been integrated with the MyProxy Service [25]. The basic computational infrastructure in neuGRID is based on gLite which provides an open source and mature Grid middleware.

The analysis environment, as shown in figure 5, enables clients to perform various functions such as the submission of workflows, tracking progress and control functions to monitor workflows. These workflows are translated by the Pipeline Service translation component and forwarded to a planner to optimize the workflow and eventually submit it to the neuGRID infrastructure for execution. For this purpose an object-oriented workflow API has been designed and implemented. The translation component implements an API which allows the translation of various workflow specification formats to a common format. The neuGRID Provenance Service is based on the CRISTAL [26] software. CRISTAL is a data and workflow tracking (i.e. provenance) system. Using the facilities for description and dynamic modification in CRISTAL in a generic and reusable manner, CRISTAL is able to provide modifiable and reconfigurable workflows. It uses the so-called description-driven nature of the CRISTAL models [27] to act dynamically on process instances already running and can thus intervene in the actual process instances during execution.

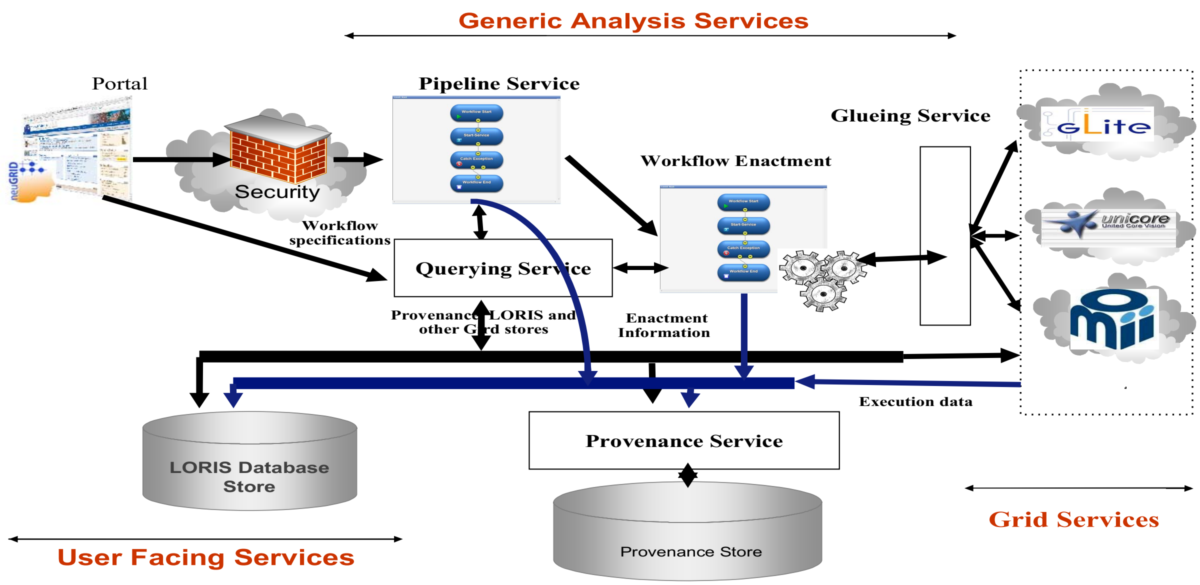

Figure 4: neuGRID Service Oriented Analysis Environment

This approach can extend and enhance the reusability of already developed services across domains and applications Users can, using the Glueing Service, 'gridify' their applications without installing and maintaining too many Grid specific libraries.The neuGRID Querying Service, apart form providing a basic interface to query data sets, provides additional reasoning and query processing facilities. It provides mechanisms to enrich and optimize user queries for efficient and rapid retrieval of the results. Communication with the underlying data resources is based on OGSA-DAI [28] which was selected since it is already widely utilised in querying heterogeneous data resources. The service provides an intermediary layer between the underlying data resources and the user. This intermediary layer is necessary to manipulate the query before it reaches the underlying data resources. This means that while using a service it remains agnostic with respect to the abstraction used to access these resources. This is also beneficial since the querying logic remains enveloped and separate within the service and not just placed within the data access abstraction. For this reason, the querying service is scalable since the querying logic placed in the service must evolve alongside newer versions of the data access abstraction.

VI. CONCLUSIONS AND FUTURE WORK

Alzheimer's disease is the most common cause of dementia, accounting for around two thirds of cases in the elderly. The study of Alzheimer's disease, its causes, its symptoms and especially its early diagnosis is now a major driver in the provision of healthcare for the elderly. By taking Alzheimer's disease as an exemplar, the neuGRID project has developed a set of analysis services and a Grid infrastructure which can enable the European neuroscience community to carry out research required for the study of degenerative brain diseases. Using the services in the neuGRID infrastructure, neuroscientists will be able to identify neurodegenerative disease markers through the analysis of 3D magnetic resonance brain images. The major benefit should be earlier diagnosis and faster development of new drugs, which will improve the quality of life of the elderly people. The set of services has been designed and developed using the service oriented architecture paradigm and the services can thereby be reusable both across Grid-based neurological data and later for wider medical analyses. The services have been developed following the SOA approach which provides the basis for extensibility and inter-operability.

Future work includes measures to demonstrate the capacity and capability of the services on different Grid middleware. Currently neuGRID is actively pursuing collaboration both with other projects in neuroscience (including in Europe NeuroLOG [10], in Canada CBRAIN [29] and in the US BIRN / LONI [18]) in the follow-on N4U (neuGRID for Users) project and with other clinical researchers in other medical domains (for example cardiology, paediatrics and cancer studies) who may wish to reuse the neuGRID services for their analyses. The design principles that have been followed in delivering a set of generalized data analysis services will provide the interoperability needed to integrate such systems.

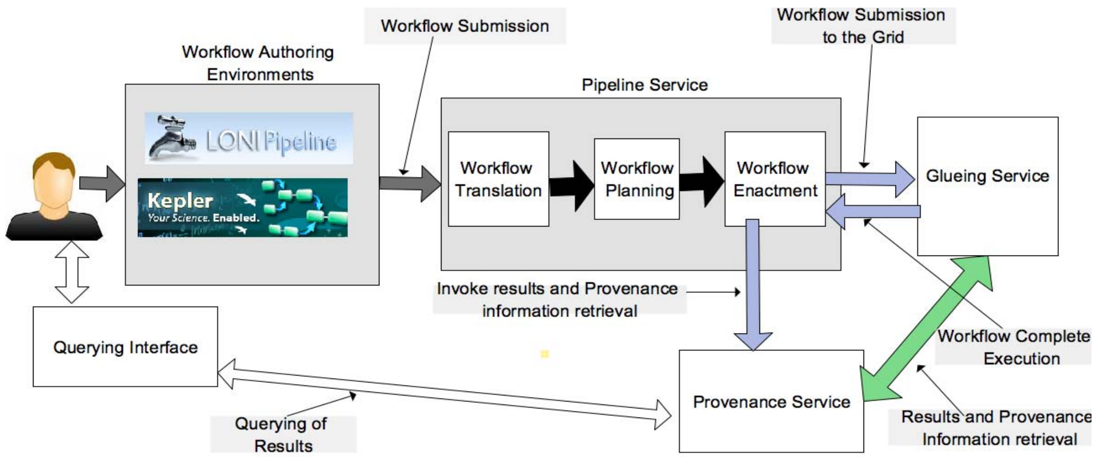

Figure 5: The Pipeline Service in neuGRID

.

REFERENCES

[1]http://alzheimers.org.uk/site/scripts/documents_info.php?documentID=936. The Alzheimer's Society..

[2] Frisoni GB et al., Neuroimaging tools to rate regional atrophy, subcortical cerebrovascular disease, and regional cerebral blood flow and metabolism: consensus paper of the EADC. J Neurol Neurosurg Psychiatry 2003;74:1371–1381

[3] Frisoni, GB et al., The clinical use of structural MRI in Alzheimer disease, Nature Reviews Neurology 2010;6:67–77

[4] Caroli A and Frisoni GB, Quantitative evaluation of Alzheimer's disease. Expert Reviews of Medical Devices 6, 569-588, 2009.

[5] Teipel SJ, Meindl T, Grinberg L, et al., Novel MRI techniques in the assessment of dementia. European Journal of Nuclear Medicine and Molecular Imaging 35, S58-S69,2008.

[6] Ashburner J, et al., Computer-assisted imaging to assess brain structure in healthy and diseased brains. The Lancet Neurology 2, 79-88, 2003.

[7] Otte A, Halsband U, Brain Imaging tools in neurosciences. Journal of Physiology 99, 281-292,2006.

[8] Frisoni GB, Structural imaging in the clinical diagnosis of Alzheimer's disease: problems and tools. Journal of Neurology, Neurosurgery & Psychiatry 70, 711-718 ,2001.

[9] Clifford RJ, Bernstein M, Fox NC, et al., The Alzheimer's Disease Neuroimaging Initiative (ADNI): MRI Methods. Journal of Magnetic Resonance Imaging 27(4), 685-691,2008.

[10] Montagnat J, Gaignard A, Lingrand D, Rojas Balderrama J, Collet P, Lahire P. NeuroLOG: a community-driven middleware design. Proc of HealthGrid'08, IOS Press, Chicago, 2008.

[11] Joseph S, NeuroGrid: Semantically Routing Queries in Peer-to-Peer Networks, In Proceedings of the International Workshop on Peer-to-Peer Computing (co-located with Networking 2002), Pisa, Italy, May 2002.

[12] Zijdenbos AP, Forghani R, Evans AC, Automatic 'pipeline' analysis of 3D MRI data for clinical trials: Application to multiple sclerosis, IEEE Trans Med Imag 21(10): 1280–1291

[13] Hauer T, et al., 2007, Experiences of engineering Grid-based medical software. International Journal of Medical Informatics. 76(8), 621-632.

[14] DICOM, Digital Imaging and Communications in Medicine, http://medical.nema.org.

[15] Kretsis A, Kokkinos P, Varvarigos E, Developing Scheduling Policies in gLite Middleware, pp.20-27, 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid, 2009, Shanghai, China.

[16] Neelin P, MacDonald D, Collins DL, Evans AC, The MINC file format: From bytes to brains, 4th International Conference on Functional Mapping of the Human Brain, Montreal, June 1998.

[17] Mildenberger P, Eichelberg M, Martin E, Introduction to the DICOM standard, European Radiology, Vol. 12, No. 4, April 2002, pp. 920-92.

[18] Pan M, Rex D, and Toga A, The LONI Pipeline Processing Environment: Improvements for Neuroimaging Analysis Research, 11th Annual Meeting of the Organization for Human Brain Mapping, 2005.

[19] Altintas I, Berkley C, Jaeger E et al., Kepler: an extensible system for design and execution of scientific workflows, Scientific and Statistical Database Management, Jan 1, 2004.

[20] Giedd, JN et al., Brain development during childhood and adolescence: A longitudinal MRI study Nature Neuroscience 2(10): 861–863), 1999.

[21] Foster I, Globus Toolkit Version 4: Software for Service-Oriented Systems, IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp 2-13, 2006.

[22] Streit A, UNICORE: Getting to the heart of Grid technologies, eStrategies | Projects, 9th edition, British Publishers Ltd, Mar. 2009, pp.8-9

[23] A Simple API for Grid Applications (SAGA) www.ogf.org/documents/GFD.90.pdf

[24] Apache Axis 2 Framework, http://ws.apache.org/axis2/

[25] Novotny J, Tuecke S, Welch V, An Online Credential Repository for the Grid: MyProxy, Proceedings of the Tenth International Symposium on High Performance Distributed Computing (HPDC-10), IEEE Press, August 2001.

[26] McClatchey R et al A Distributed Workflow & Product Data Management Application for the Construction of Large Scale Scientific Apparatus, NATO ASI Series F : Computer & Systems Sciences Vol 164 pp 18-34. Springer Verlag 1998. Verlag 1998.

[27] Estrella F et al., 2003, Pattern Reification as the Basis for Description-Driven Systems Journal of Software and System Modelling Vol 2 No 2, pp 108-119 Springer-Verlag, ISSN: 1619-1366.

[28] Comito, C, Gounaris A, Sakellariou R, Talia, D, A service-oriented system for distributed data querying and integration on Grids, Future Generation Computer Systems, Volume 25, Issue 5, May 2009, 511-524.

[29] CBRAIN Project, http://cbrain.mcgill.ca/