
Glueing grids and clouds together: a service-oriented approach

Ashiq Anjum* and Richard Hill

School of Computing & Mathematics,
University of Derby,
Derby, UK
Email: a.anjum@derby.ac.uk
Email: r.hill@derby.ac.uk
*Corresponding author

Richard McClatchey

CCCS Research Centre,
FET Faculty,
UWE Bristol, UK
Email: richard.mcclatchey@cern.ch

Nik Bessis

School of Computing & Mathematics,
University of Derby,
Derby, UK
Email: n.bessis@derby.ac.uk

Andrew Branson

CCCS Research Centre,
FET Faculty,
UWE Bristol, UK
Email: andrew.branson@cern.ch

Abstract: Scientific communities are actively developing services to exploit the capabilities of service-oriented distributed systems. This exploitation requires services to be specified and developed for a range of activities such as management and scheduling of workflows and provenance capture and management. Most of these services are designed and developed for a particular community of scientific users. The constraints imposed by architectures, interfaces or platforms can restrict or even prohibit the free interchange of services between disparate scientific communities. Using the notion of 'Platform as a Service' (PaaS), we propose an architectural approach that addresses these limitations so that users can make use of a wider range of services without being concerned about the development of cross-platform middleware, wrappers or any need for bespoke applications. The proposed architecture shields the details of heterogeneous Grid/Cloud infrastructure within a brokering environment, thus enabling users to concentrate on the specification of higher level services.

A. Anjum et al.

Keywords: grid and cloud computing; services oriented architecture; platform as a service.

Reference to this paper should be made as follows: Anjum, A., Hill, R., McClatchey, R., Bessis, N. and Branson, A. (XXXX) 'Glueing grids and clouds together: a service-oriented approach', *Int. J. Web and Grid Services*, Vol. X, No. Y, pp.xxx–xxx.

Biographical notes: Ashiq Anjum is a Lecturer in Computer Science at University of Derby (formerly at University of the West of England, Bristol UWE and Imperial College London), and is carrying out research work in distributed systems, grid/cloud computing and data mining. He is also contributing to various European and US research projects and is an active researcher in the area of Cloud/Grid and distributed systems. He has authored more than 50 peer-reviewed publications.

Richard Hill is Head of Subject, Computing and Mathematics, and a Professor of Intelligent Systems at the University of Derby. His research interests include agent oriented software engineering, collective intelligence, emergence in systems, and enterprise cloud computing. He has published over 100 articles and edited several books and journal special issues.

Richard McClatchey has been research active for the past 30 years in the areas of large-scale distributed data and process management and database. Currently a Fellow of both the British Computer Society and the Institute of Engineering and Technology with a PhD in Physics and DPhil in Computer Science, he has published over 150 papers and has holds the Chair of Applied Computer Science at UWE. His current expertise lies in grid data and knowledge management and particularly in its application to medical applications.

Nik Bessis is a Professor and Head of the Distributed and Intelligent Systems (DISYS) research group in the School of Computing and Mathematics at University of Derby, UK. His research interests focus on resource management, data integration, annotation and data push methods and services in distributed dynamic environments such as grids, clouds, interclouds and IoT. He is involved in and leading a number of projects. He has published over 150 papers, won 2 best paper awards and is an editor of several books and the Editor-in-Chief of the *IJDST*. He serves as a conferences chair and keynotes speaker.

Andrew Branson is a Research fellow in the Centre for Complex Cooperative Systems at UWE, Bristol. He is the principal software architect of the CRISTAL software originally produced at UWE and CERN to manage the construction of the Compact Muon Solenoid experiment at the Large Hadron Collider (LHC). He has worked on a number of EC projects including Health-e-Child and the FP7 projects neuGRID and N4U.

This paper is a revised and expanded version of a paper entitled 'Gluing grids and clouds together: a service oriented approach' presented at the '2nd International Conference on Emerging Intelligent Data and Web Technologies', Polytechnic University of Tirana, Tirana, Albania, 7–9 September 2011.

1 Introduction

Scientific communities are increasingly building high-level distributed services such as analysis services, workflow and scheduling services, monitoring services and provenance services. These services or applications help scientific users to analyse their data, to manage, to monitor and optimise computational infrastructure, to extract knowledge and to share it with other users. Most of these services are designed and developed for a particular computing platform [such as the gLite Grid middleware; Kretsis et al. (2009) or the OpenStack Cloud platform; OpenStack (2011)] and for a specific community of scientific users, for example molecular or biomedical analysis (Koehler et al., 2010). It is often difficult to share or re-use services from one community with another community due to architecture, interface or platform limitations. Similarly, the services developed for one platform such as Microsoft Azure (Microsoft Windows Azure Platform, 2011) cannot be deployed in another computing environment such as Amazon Web Services (2011a) due to stack mismatch and interoperability issues between the service providers. This has become particularly critical in the case of distributed infrastructures such as Grids and Clouds that host a number of high-level applications.

An application designed and developed for one Grid middleware, despite resource sharing, is often challenging to run on another platform thus limiting the wide scale adoption of this technology. This is particularly prevalent for Cloud providers, since there is a tendency to adopt proprietary standards for service implementation, which inhibits the generalisation of service specifications, and leads to a consequent Cloud platform lock-in. As Grid technologies evolve onto Cloud platforms the situation worsens further since either the applications need to be redeveloped for the emerging platforms or significant changes are required to adapt the applications for the newer platforms. This is further compounded when Cloud stacks offer almost similar features, and services cannot cooperate due to technical or architectural limitations. Choosing a Cloud infrastructure or platform means a commitment for the whole application execution life. As such, an application is bound to a platform and users have limited choice if they intend to change an infrastructure provider in future. The potential wasted effort as well as the added investment required presents a significant barrier to the adoption of Cloud technologies. Since applications for Grids and Clouds are relatively tightly coupled to the underlying platform and we are presented with two significant issues. Firstly, typical Grid and Cloud services such as resource management, scheduling, monitoring and security policy management are specific to the platform and therefore users must acquire and maintain specific expertise related to the Grid or Cloud infrastructure used. Secondly, an application that itself exhibits heterogeneity before deployment to a Grid or Cloud requires a specific implementation that further increases the coupling of the application to a particular infrastructure. Grid and Cloud infrastructures and platforms are designed in such a way that an application has to be horizontally as well as vertically integrated with the underlying platform to achieve the desired performance guarantees, throughput and optimisation. Isolation, encapsulation, ad hoc communication and transparent contract negotiation between the platform and applications has generally not been considered in the design and development of the Cloud as well as Grid platforms.

As a result of these design practices, any changes in platform middleware, or a desire to move from one Grid or Cloud provider to another, will require a significant amount of application re-development (Uehara, 2011). Consequently users have a fundamental dilemma in terms of weighing-up the eventual value of the system against fundamental application re-development costs. Though most Grid and Cloud platforms claim to

follow a service-oriented paradigm, the current stacks only partially implement Service-Oriented Architectures (SOA), and most of the time any benefits of SOA cannot be exploited due to the architectural choices that have been followed to implement the Cloud and Grid platforms. This clearly does not enable the benefits of a service-oriented approach to be realised, whereby users should be concentrating upon their core business, rather than the engineering requirements of their underlying infrastructure.

Whilst the marketing hype surrounding Grid and Cloud computing often promises elasticity of resources, ubiquitous access and abstraction from heterogeneity (Jeffery and Neidecker-Lutz, 2010), the reality is often somewhat different, and as long as application portability, compatibility and interoperability with underlying infrastructures remain a challenge (Kertész and Kacsuk, 2009), the true potential of the Grid and Clouds platforms cannot be realised. For instance, if an application's needs cannot be addressed by a local Grid/Cloud infrastructure, the ability to add capability on demand by calling upon third party Grid or Cloud services, irrespective of where and how they are hosted, is still problematic and far from the promise of 'utility computing'. As described earlier, service-based infrastructures tend to be constrained by the suppliers of the platforms, thus acting as a disincentive towards the wider utilisation of Grid and Cloud platforms. If an application is built upon a combination of local resources and commercial resources (such as Amazon Web Services, 2011b; Rack Space Cloud Hosting, 2011), any future or emerging business requirement to adapt or transfer to other platforms will require application re-development, which does not faithfully follow a service-based model of computing. A purer form of Platform as a Service (PaaS) and IaaS should enable the users to run their applications as a utility consumption, without being concerned with whoever is providing the infrastructure, what platform or protocol has been used or how to switch to a different computing infrastructure provider. The non-availability of such an environment is therefore a fundamental barrier towards the pervasive adoption of Grid and Cloud platforms demanding concrete steps to be taken to address the limitations.

This is far from ideal for users, who not only have preferences in using tools and technologies for their analysis, but also may resist the introduction of new interfaces or attempts to customise their applications. This takes up significant time and resources in understanding the platforms and technologies, which should be spent on solving the scientific or business challenge. Essentially, this trend drives users away from their core business and discourages them from exploiting Grid or Cloud resources. If a decision is made to exploit these alternative platforms, then users must invest considerable time to understand the platform, its interfaces, underlying protocols and technologies before high-level applications to exploit the computer and data resources available can be written. There is no platform or service available that can enable users to abstract these details so that they only focus on the application level, with the underlying complex infrastructure details remaining hidden (unless, of course, it is important for increasing the performance of an application or for other users' needs). The problem is exacerbated by frequently changing platforms, technologies and interfaces, which force users to re-design and re-implement their applications. This restricts the adoption of computing platforms made available through Grid and Cloud technologies, which are anticipated as deployment platforms for better performance, scalability and serviceability of applications.

In this paper, we propose a platform that addresses the limitations stated above. We extend the notion of PaaS to address platform, interface and technology limitations, to ensure that users do not need to worry about the lower level details of a platform for cross-community applications. Using a service-oriented platform, a brokering environment is offered which shields the heterogeneity of distributed resources present in Grids and

Glueing grids and clouds together

Clouds. This platform (referred to as the Glueing Platform (GP) thereafter) supports applications by allowing concurrent access to a number of platforms, by shielding the underlying infrastructure details from users and applications, thus enabling users to implement higher level services. The platform provides features such as workflow enactment, execution control and monitoring, provenance and tracking, concurrency management at the application level, scalability to support a number of applications and users, and security. The GP exposes the Cloud Data Management Interface (CDMI) (SNIA, 2011), Open Cloud Computing Interface (OCCI, 2011; OGF Working Group, 2011), Open Virtualization Format (OVF) (Wikipedia, 2011) and Simple API for Grid Applications (SAGA) (Jha et al., 2007) implementations as a distributed service using the service-oriented paradigm. Briefly, the key advantages of the GP are:

- It provides (using OCCI and SAGA) a standard way of accessing Cloud and Grid services without tying down services and applications to a particular Cloud or Grid middleware.
- It enables access to any deployed Grid or Cloud middleware through a service-oriented brokering environment.
- It offers a solution that enhances re-usability of existing services across disparate domains and applications.
- It facilitates development productivity by offering a service-based approach to shield users from complex Grid and Cloud functionality.
- It offers an easy to use approach for enabling clients to port their applications to Grids or Clouds (or both), without installing and maintaining too many Grid and Cloud specific libraries.

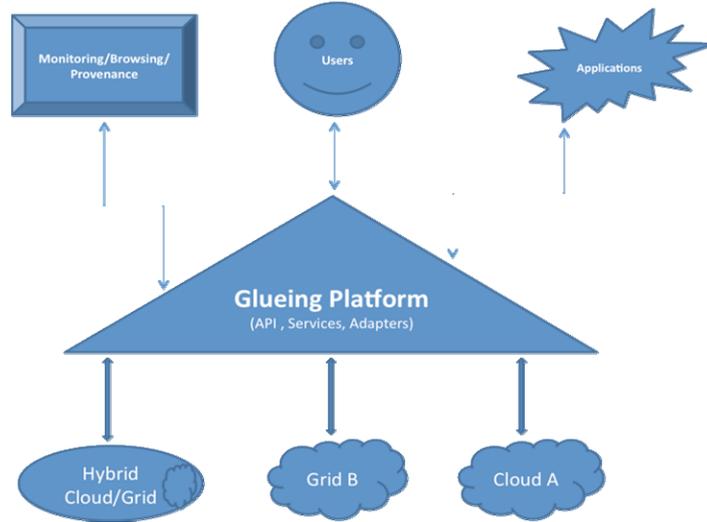
The paper proceeds as follows. Section 2 discusses the GP architecture and describes the considerations that went into the design of the platform. Section 3 presents the role of Cloud standards in the Glueing Platform and Section 4 describes an exemplar case study about the usage of the platform. Section 5 describes the enactment and execution details of the platform. The services and components in the platform are then described in Section 6. These are the core APIs that enable the user to access the platform functionality. Section 7 concludes the paper and identifies opportunities for future work.

2 The Glueing platform architecture

The high-level architecture for the Glueing Platform (GP) is shown in Figure 1. The GP sits between the infrastructure and the applications and provides the required independence from the underlying platforms. A user does not need to know about the lower level APIs or protocols, as components and services are provided to enable access to the underlying compute and data resources in a Grid or Cloud. The services in the platform include authentication and authorisation services, scheduling and resource management services, monitoring and provenance services, data access and browsing services and enactment and execution services. To interact with heterogeneous middleware, an adapter-based approach has been followed, that implements middleware for each supported Grid or Cloud in the GP. The GP enables applications to be platform agnostic by exposing suitable interfaces that can be accessed to run an application, to browse data, to monitor the execution details or to retrieve the provenance history. The user side

interfaces remain the same irrespective of the underlying platform or infrastructure and the GP takes care of the evolving interfaces and components at the platform and infrastructure levels. The API details are described in the following sections.

Figure 1 The Glueing Platform architecture (see online version for colours)



The GP exposes the OCCI as a web service in order to achieve interoperability and portability across the Cloud platforms. The OCCI interface provides an API specification for remote management, monitoring, scaling and deployment of Cloud computing infrastructure. All aspects of virtualised service life cycle management have been considered in the OCCI specification and its implementations. The OCCI specification depends on the CDMI standard, developed by Storage Network Industry Association (SNIA), to create, update, retrieve and delete data from Cloud. The CDMI manages all aspects of data movement and migration between storage services in Clouds through standard interfaces. There are some overlaps between CDMI and OCCI interfaces but in general the consensus is that data and storage operations come under CDMI and compute, security and infrastructure issues are dealt with OCCI. Further details on how OCCI and CDMI are integrated and how data and virtual machines are exported using OVF are discussed in Section 3.

The GP exposes the SAGA implementation as a web service to address the objectives described previously. SAGA is an open standard that is defined and maintained by the Open Grid Forum (OGF, 2011), which describes a high-level interface for easy programming of Grid/Cloud applications. SAGA is not designed for middleware developers; rather it provides APIs for developing applications using different Grid or Cloud middleware. The SAGA API permits a job to be executed on a particular middleware by the runtime loading of a middleware adaptor. SAGA middleware adaptors pass jobs to the middleware as its clients and are responsible for low-level communication with it.

The high-level interfaces, provided by SAGA, for communicating with different middlewares, remove the need for the use of middleware specific APIs. For example, if an application intends to use three different middlewares, then the application has to use a different API set for interacting with each middleware. SAGA, on the other hand,

Glueing grids and clouds together

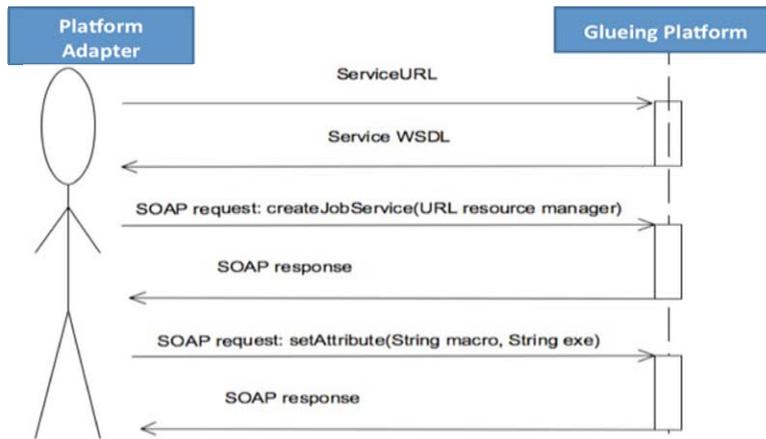
provides a single generic API to interact with every middleware whose adaptor is loaded in the application. Thus, the GP, using SAGA, 'glues' a number of client applications together with different middleware using the relevant adaptors. The GP wraps the SAGA API and provides a WSDL binding for client applications. There are numerous advantages of this approach, compared to more traditional means whereby the client applications directly use the SAGA API implementation. The most important reasons for adapting a web service-based approach to SAGA are now detailed.

SAGA implementations come packaged in two parts: middleware adaptors and the API. The client uses the API to communicate to the Grid middleware via the appropriate adaptor. In order for the adaptor to communicate with the Grid middleware the client application, SAGA API and the adaptor must be deployed on a system which has the core middleware packages installed, deployed and configured. For instance, consider a user who wants to submit a job to a Condor (Thain et al., 2005) cluster deployed as part of an OMII (Distributed Systems Technology Centre, 2004) middleware Grid. The client application must be deployed on a node, which has the OMII middleware client installed, and Condor installed and configured in submit only mode. If in future the application is ported to another middleware, the client application must be deployed on a machine which has the new middleware installed. This increases porting costs and introduces complex installation and deployment procedures for applications which make the life of common users quite difficult. SAGA is an evolving standard. Currently, there are numerous things which are not provided by the SAGA API. Workflows are one of them and discovery is another. The Glueing Platform can be used as a means to providing functionality, which is not currently provided by SAGA. In this scenario, the Glueing Platform would expose all of SAGA's capabilities. In addition it would also expose higher level functions which SAGA does not provide; however, they deal with the Grid middleware and are required by client services.

The GP shields users from these highlighted difficulties, and deploys the adaptors and middleware at the location of the GP host. Clients interact with the GP via a standard web service-based infrastructure. In this scenario, if the middleware needs to be changed, all that the clients require a new end point to an appropriately deployed GP. In order to make the GP SAGA compliant, two separate WSDL descriptions have been provided, one of which points to the full SAGA implementation, and the other describes the extended functionality (which is not currently supported by SAGA). Both WSDL descriptions map to the GP. Additionally, the GP exposes SAGA API functions as web service methods and provides a one-to-one correspondence to the SAGA API functions. Client applications can transparently access the GP by using a SAGA SOAP adaptor, which is an implementation of the adaptor API provided by SAGA. The clients can include the SAGA SOAP adaptor and can write applications using the standard SAGA API classes. The SAGA API calls, generated on the client side, are passed to the GP by the SAGA SOAP adaptor, which is responsible for all communication with the GP. The SAGA SOAP adaptor is a middleware between the client applications and the GP. The client applications define, create and submit jobs according to the standard specification of SAGA (Goodale et al., 2008). These instructions are then translated into SOAP requests by the SAGA SOAP adaptor. SOAP requests are used for communication with the Glueing Platform. The Glueing Platform actually executes the client instructions using SAGA APIs and middleware adaptors. The SAGA SOAP adaptor has methods to discover the GP and to retrieve its end point URL. The end point URL is used to access the service WSDL, which is then used for service invocation. The WSDL describes the definition of all the exposed methods. The SAGA SOAP adaptor calls the published methods using SOAP requests and the Glueing Platform sends back the SOAP response

to the SAGA SOAP adaptor. The SAGA SOAP adaptor then translates the SOAP response and returns the execution results to the client in the form of Java or SAGA specific objects. Figure 2 shows a scenario where a SAGA SOAP adaptor interacts with the GP. The SAGA SOAP adaptor passes the middleware and job information to the Glueing Platform using SOAP objects and the SOAP response is sent back to the SAGA SOAP adaptor from the GP.

Figure 2 A sample use case (see online version for colours)



3 Cloud standards and the Glueing platform

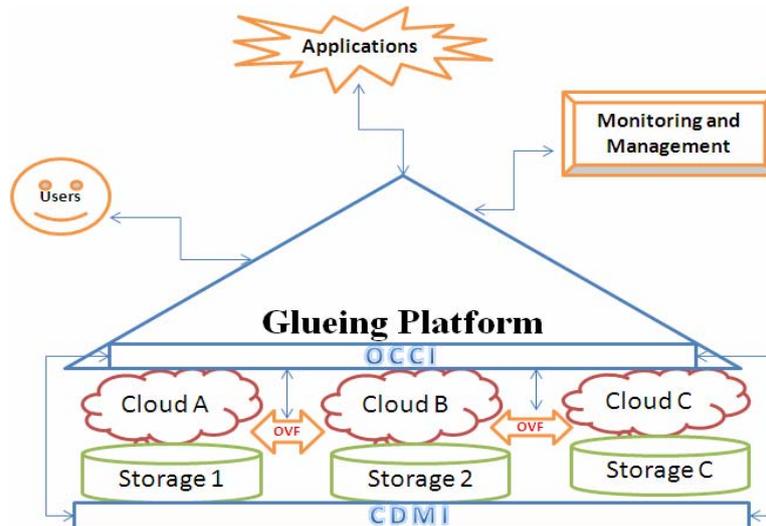
As stated earlier, the standard specifications have been developed to achieve interoperability, portability and to provide elegant communication and integration between heterogeneous components and services that access various features from different Cloud platforms. Cloud platforms are proprietary and vendors have implemented their own interfaces and access mechanisms, and without standard interfaces and specifications, it will be impossible to achieve infrastructure interoperability and application portability across various Cloud platforms. Sensing the gravity of the situation, organisations have worked together to develop industry and community standards for seamless communication, integration and interoperability of Cloud stacks and platforms. The three most important issues that have been addressed include the following:

- OCCI: monitoring, management, scalability and deployment of Cloud infrastructure. This allows runtime management of compute infrastructure.
- CDMI: creation, retrieval, removal and updating of data sets. This allows runtime management of storage infrastructures.
- OVF: exchange of data and virtual machines in a standard machine readable format. This provides a means to package virtual infrastructures that can be imported or exported based on the service deployment requirements.

Figure 3 shows the communication that happens between the Glueing Platform and these standard implementations to achieve application portability and infrastructure interoperability.

Glueing grids and clouds together

Figure 3 Clouds standards and the Glueing Platform (see online version for colours)



For example, a user may plan to deploy a Cloud service for a scientific analysis. The first thing the user should do is to create the service, test it and provide the necessary resources to make the service scalable so that it can cope with the user demand. Once the basic building blocks have been made available and the scalability issues have been resolved, the next important thing is to make the deployment exportable so that the same deployment can be made available in case of a failure or if more than one sites need to be commissioned for fault tolerance or another infrastructure providers need to be considered for cost savings.

The service deployment, availability management and scalability are addressed by the OCCI. The Glueing Platform will trigger the OCCI interface to deploy the service and carry out the scalability and management operations. The storage management is a function of the CDMI and the Glueing Platform invokes the CDMI interface to store, retrieve, update and remove files from the Cloud storage devices. In case the service was already deployed and a new instance of the service needs to be created, the Glueing Platform will call the OVF interface to export or import the service instances. The virtual machine instances can be migrated between the resources within or between the data centres using the Distributed Management Task Force (DMTF) based specification and implementations of the OVF standard. The Glueing Platform, using the standards specifications and tools, offers platform and infrastructure portability. Users are free to move their applications and services across different Cloud platforms and infrastructures.

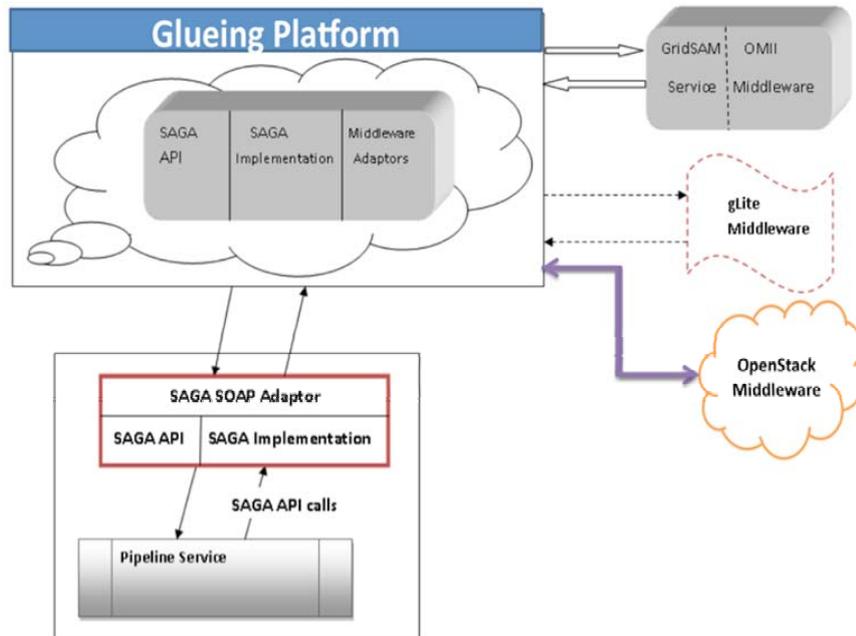
The Glueing Platform is the first point of contact for users and applications to deploy or use services and to manage and monitor the infrastructures. The applications do not have to access platform specific libraries or components to provide the service and are free to move to another infrastructure or a Cloud platform. The Glueing Platform takes care of all the low-level details that deal with the standards interfaces and components. The Glueing Platform architecture is flexible enough to accommodate new requirements and standards such as audit and security from Cloud Security Alliance (CSA). The Glueing Platform brings all these into a single framework and users are provided a

simple to access service and in return they get benefits from all the standard-based implementations of tools and technologies for deploying, operating and managing the infrastructures. This approach makes the applications and services to be widely deployed on a number of platforms, avoids vendor lock-ins and promotes interoperability and application portability.

4 Case study

The architecture diagram (Figure 4) shows how an application contacts the GP. The pipeline service, as shown in Figure 4, is one of the potential applications of the GP. This service uses the SAGA APIs and SAGA SOAP adaptors to communicate with the GP. The SAGA SOAP adaptor accesses different methods of the GP by getting its WSDL. The methods published in the WSDL execute the actual instructions which are generated on the client side. Thus, the pipeline service initiates a Grid activity which is then forwarded to the GP by the SAGA SOAP adaptor. The GP, as shown in Figure 4, can communicate with different Grid middleware such as OMII/GridSAM (Lee et al., 2005) or gLite, etc. This allows client applications, such as pipeline service, to use Grid resources provided through different middleware. The GP uses the JavaGAT (van Nieuwpoort et al., 2007) middleware, which is a SAGA implementation, and OMII adapters to communicate with the OMII middleware.

Figure 4 Glueing Platform and grid middleware (see online version for colours)



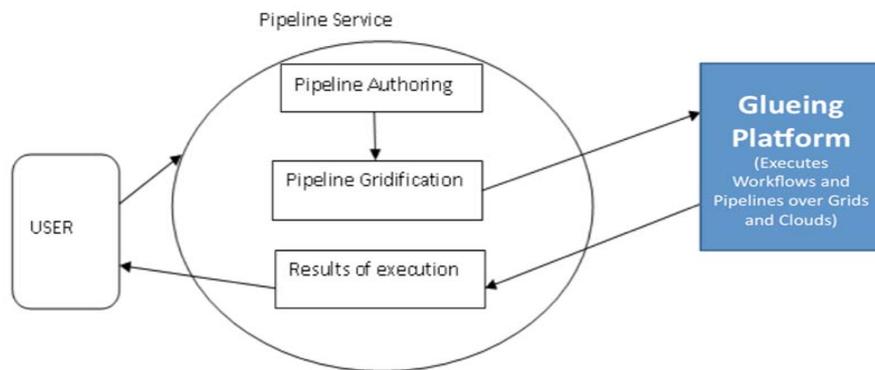
GP client applications, such as the pipeline service in the neuGrid project (Redolfi et al., 2009), allow clinicians and neuroscientists to create neuro-imaging pipelines, in a user friendly environment, for a series of automated transformations on brain images. Once

Glueing grids and clouds together

the pipelines are constructed, the neuro-imaging pipeline service parallelises them. The processes defined in the pipelines are usually very compute-intensive and deal with large amounts of data. Therefore, in the next step the workflows are modified for their optimal execution over the Grid.

The neuro-imaging pipeline service, using the GP, submits the workflows to the Grid for their execution and the results of execution are returned to the client interface of pipeline service. A simplified design of the pipeline service is shown in Figure 5.

Figure 5 Pipeline service (see online version for colours)



The neuro-imaging pipeline service passes workflows, or sequences of processes, to the Glueing Platform. The Glueing Platform exploits the SAGA system and executes the processes in the workflows on the Grid using appropriate Grid middleware. The results of execution are passed over to a provenance service and eventually reach back to the client.

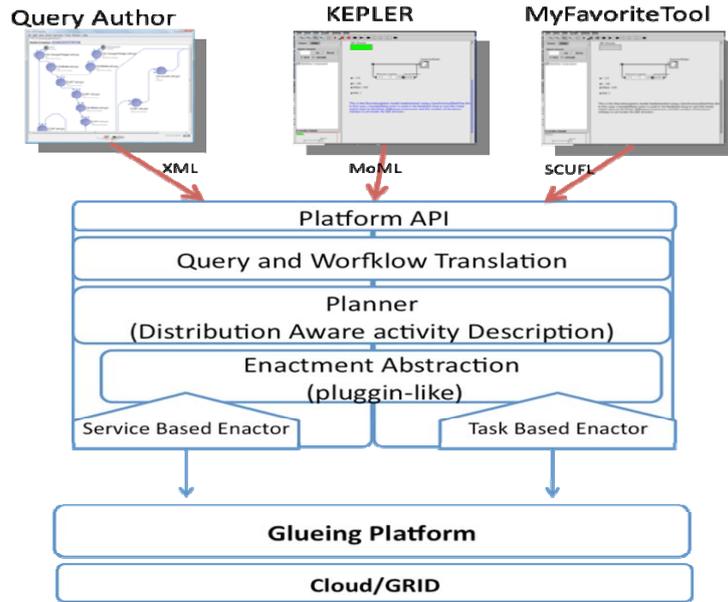
5 Execution and enactment in the Glueing platform

The Glueing Platform offers an environment where workflows and applications can be executed in a platform neutral manner. The components of the Glueing Platform to achieve a platform-agnostic enactment are shown in Figure 6. The Platform treats every user activity, whether it may be an application, user query or series of actions, as a workflow. This enables the platform to perform a series of actions on the queries or users' steps to plan and orchestrate the execution for a maximum throughput. This also enables the users to optimise their queries and workflows by performing a number of intelligent actions based on the concepts described by Cho et al. (2011) during the planning phase to save costs and times.

Every user query or job is charged by the Cloud providers; minimising the number of jobs will save costs. Similarly, by planning their activities, users can remove redundant tasks, group similar tasks or schedule their activities on resources in Clouds which can provide users improved throughput. Considering the consequence of these actions, queries or workflows can do more in less time, and can lead to intelligent and adaptive practices in Clouds (Sun et al., 2011). Users can select resources on which they can execute their workflows or jobs and they also have control on the choice of the authoring

environment. They can author their queries or workflows in an environment such as Kepler, or they may use a simple command line tool, as their actions can be passed on to the planning component in a structured format. As shown in Figure 6, the architecture of the Glueing Platform is flexible and any suitable authoring environment can be accommodated.

Figure 6 Enactment in the Glueing Platform (see online version for colours)



After authoring a query or workflow, the users invoke their submission. At this stage, several things can happen: firstly, the authored query, process or workflow that is represented in a format which is native to the authoring environment can be translated into a common platform model. This model can be an object-oriented model of the query or workflow or this can be translated into another format which is convenient to the planning engine. After workflow translation, architecturally the workflow should be planned for its execution. In the planning stage, a number of intelligent strategies can be evaluated for an optimal execution of the workflows or processes. The planned workflows or jobs can then be submitted to a Cloud or Grid for execution. The execution can be task based or there could be pre-deployed service instances to support the execution. In both cases, an appropriate enactor prepares the environment for execution by discovering suitable resources and scheduling the jobs/workflow onto these resources by performing matchmaking and mapping processes. Once appropriate resources have been shortlisted; the enactor component schedules the jobs/workflows for execution. A monitoring component keeps track of the execution process and a status update is reported to the user in intervals. The results, execution traces and their provenance logs are sent to the users to make it a transparent execution.

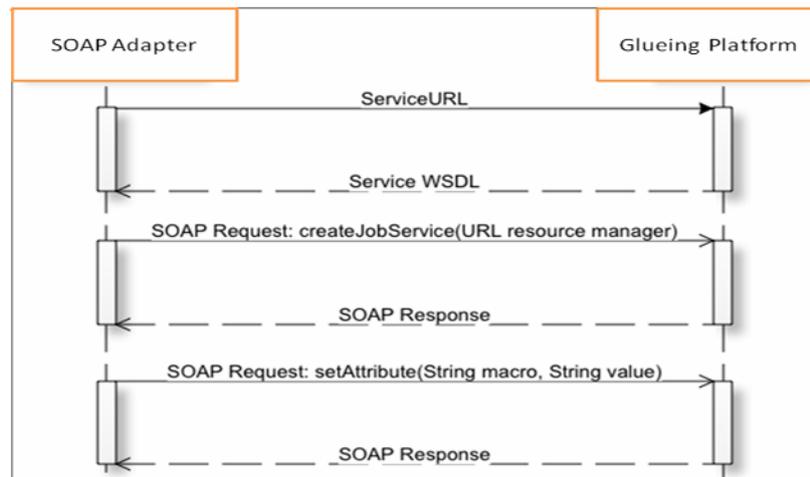
6 Platform components and functionality

The SOAP adaptor as described in Section 5 requires a service end point URL to communicate with the Glueing Platform. The end point URL is used to access the service WSDL, which is then used for service invocation. The WSDL describes the definition of all the exposed methods. The SOAP adaptor calls the published methods using SOAP requests and the Glueing Platform sends back SOAP responses to the SOAP adaptor. The SOAP adaptor then translates the SOAP response and returns the execution results to the client in the form of Java or SAGA specific objects.

Figure 7 shows a scenario where the SOAP adaptor interacts with the Glueing Platform. The SOAP adaptor passes the middleware and job information to the Glueing Platform using SOAP objects and the SOAP response is sent back to the SOAP adaptor from the Glueing Platform. The exposed functions of the Glueing Platform are specifically discussed in detail in the following sections.

The GP exposes SAGA APIs in the form of a web service (or web services) which allows client applications to create, run and monitor jobs transparently from the underlying Grid middleware. The client applications, using the SAGA SOAP adaptor, pass job and middleware information to the Glueing Platform. The platform loads the appropriate middleware adaptor, based on the middleware information. For example, if a client wants to run jobs or access data on an OMII-based infrastructure, the client specifies an OMII/GridSAM middleware; the GP sets a system property “JobService.adaptor.name” equal to “javaGAT”. This system property behaves like an environment variable for the SAGA engine to narrow down the selection of middleware adaptors. The GP then makes sure that JavaGAT loads a Grid SAM adaptor by creating an appropriate session and context for SAGA. The context is a specific piece of information that is shared with a particular session. An application may associate different contexts with a particular session in order to perform different functions in that session.

Figure 7 Information flow between SOAPAdapter and the Glueing Platform (see online version for colours)



The GP creates a preferences context for JavaGAT to make sure that it loads the GridSAM adaptor. This is done by setting the context property “ResourceBroker.adaptor.name” equal to “Gridsam”. This context is then added to the application session for further reference during the lifetime of the application. Once the session is created the GP creates a job, based on the job information sent by the client application. The GP uses SAGA job management APIs for creating a job, submitting it to the available Grid resources and retrieving its status. The SAGA job management API covers four classes: JobService, JobDescription, Job and JobSelf. The order in which the GP uses this job management and other APIs is as follows:

6.1 Selecting a resource manager

The JobService API is used to select a resource manager. An instance of JobService represents a resource manager backend. A resource manager is an end point where the job is submitted by the client application. This resource manager can also be an execution service if it executes the job.

Input parameter: To create an instance of JobService class an end point URL of resource manager is required as input parameter to create JobService method.

Example: The GP uses an end point URL for submitting the job to resource manager. For example, if resource manager is Grid SAM then an instance of JobService is created as: `<JobService> js = JobFactory.createJobService (new URL (“https://HOST:PORT/Gridsam/services/Gridsam”)).`

6.2 Job definition

The JobDescription API defines the job using a well-defined set of attributes such as the application executable and associate arguments. The JobDefinition attributes behave like tags in JSDL/JDL and thus these attributes mimic JSDL for the middleware and are passed to it internally by SAGA.

Input parameters: The JobDescription API needs two essential parameters in order to define the job, i.e. the application executable path and the exact application parameters.

Example: If the application execution is “/bin/echo” which takes a string as input parameter i.e. “hello” then the GP, using JobDescription, defines a job as: `<JobDescription> jd.setAttribute (JobDescription.EXECUTABLE, “/bin/echo”) and <JobDescription> jd.setVectorAttribute (JobDescription.ARGUMENTS, new String[] {“hello”}).`

6.3 Data stage-in/-out

The job execution results can be stored in an output file as: `<JobDescription> jd.setAttribute (JobDescription.OUTPUT, “outputFile”)`. The JobDescription class also defines resources for data staging. For example, the results of execution can be staged out at a remote location using FTP as: `<JobDescription> jd.setAttribute (JobDescription.FILETRANSFER, “file://HOST:PORT/outputFile < outputFile”)`. The execution results are staged, once execution is complete. The results can be copied by calling the copy function of the GP and by passing its source and target URL parameters.

Glueing grids and clouds together

6.4 Job creation

The GP creates an instance of the Job class using createJob function of JobService, which takes an instance of the JobDescription class as input parameter. Instances of both JobService and JobDescription classes are a pre-requisite for creating a job.

Example: A job is created as: `<Job> j = <JobService> js.createJob (<JobDescription> jd)`. The GP executes jobs by `<Job>j.run()` and an implementation of Callback interface allows it to get monitoring information of job during the course of execution.

6.5 Asynchronous job execution

The GP also allows the client applications to run different independent jobs at the same time. The feature of running jobs asynchronously is provided by the SAGA TaskContainer API. An instance of TaskContainer may contain multiple tasks. A task is a SAGA API call, whereas a job is a remotely running application. The SAGA job class extends task interface; therefore, a TaskContainer can also contain multiple jobs. The jobs are executed asynchronously when the instance of TaskContainer calls its run() method.

Example: If there are two different jobs, j1 and j2 and both copies a file from one URL to another then the TaskContainer will add both jobs as: `<TaskContainer> T.add(j1); <TaskContainer> T.add(j2)`. The task container will execute these jobs asynchronously as: `T.run()`.

6.6 Replica management

The GP also allows client applications to replicate files over Grid resources using FTP or GridFTP. If client wants to replicate a file to a Grid resource using GridFTP, he/she can do this by calling replicate function of the GP by providing it the location of GridFTP server and the source file.

Example: If the location of the source file is `/home/test.txt`, the replica is generated by first creating an instance of for the source file as: `LogicalFile lf = new LogicalFile (“/home/test.txt”)` and then this file is replicated to a remote server at the server URL as: `lf.replicate (new URL(“ftp://server:port/test.txt”))`.

6.7 Job monitoring

The GP monitors a job during its lifetime using callbacks, and it provides two methods of retrieving monitoring information. One is to use the higher level implementation, by which the client application can invoke the monitor (String jobId) function of the GP to get complete monitoring information, associated with a particular job id. In this implementation, the GP stores the monitoring information of the job in a string and it is returned to the client when (s)he invokes the monitor. The other method of monitoring a job is to use a SAGA API function. In this implementation the client first gets the job metric and it can then be used to get the metric name and value.

Example: A job metric can be obtained by using a job instance and the metricName. If j is an instance of job as created in Section 6.4 and metricName is `Job.JOB_MEMORYUSE` then the job metric can be obtained as: `Metric m = getMetric (j, Job.JOB_MEMORYUSE)`. This job metric contains the information of memory used by the particular job j, while it is in execution. The metric value can be obtained by `m.value`.

6.8 Data querying

The GP also provides a mechanism for querying the data, staged out or replicated at the remote locations. The staging end points may contain a number of execution results in the form of files. A client may want to query a particular pattern to see the list of files present at the remote location. This can be achieved by calling the find(NSDirectory nsDir, String pattern) function, exposed by the GP.

Example: The GP queries a particular pattern using the instance of NSDirectory. The instance of NSDirectory points to a remote directory. If the client queries a pattern “*.jpg” to list all the jpg files, staged out at the remote directory then the GP fetches the list as: nsDir.find(“*.jpg”).

7 Conclusions and future directions

The Glueing Platform addresses some major requirements of users who wish to adopt a flexible approach to accessing the Grid and Cloud. The heterogeneity of distributed resources and details of Grid or Cloud middleware architectures will be obscured from users. The GP also hides complexities of interfacing with different Grid and Cloud middleware, which will allow access to Grid and Cloud resources through a set of high-level functions. This is achieved by exposing CDMI, OCCI, OVF and SAGA APIs, all communication being achieved through middleware adaptors. This shields the low-level middleware difficulties from the user and assists them in using resources with little knowledge. The design of the GP is based on service-oriented architecture principles, which will help different services and applications to use service functionalities through standardised interfaces. This will also allow other client applications to use the Platform features by inspecting its WSDL, available online at the service end point URL.

Future work includes the provision of support for resource discovery [building on the work reported by Chuang and Kao (2011)] and improving it to a level so that it may be used for production quality needs. Support for workflow composition and enactment is another major issue that needs to be addressed (Hasham et al., 2011). The support for Cloud platforms such as OpenStack and EUCALYPTUS (Nurmi et al., 2009) also needs to be provided, along with support for application hosting, job distribution and concurrency. The infrastructure interoperability across disparate platforms also needs a solution for semantic interfaces in Clouds and Grids (Flahive et al., 2009; Flahive et al., 2011) and this is one of the services that we need to integrate in the Glueing Platform in future.

References

- Amazon Web Services (AWS) (2011a) *Innovation*. Available online at: <http://aws.amazon.com/> (accessed on 5 December 2011).
- Amazon Web Services (AWS) (2011b) *Amazon Elastic Compute Cloud*. Available online at: <http://aws.amazon.com/ec2/> (accessed on 5 December 2011).
- Cho, Y., Moon, J., Kim, I., Choi, J. and Yoe, H. (2011) ‘Towards a smart service based on a context-aware workflow model in u-agriculture’, *International Journal of Web and Grid Services*, Vol. 7, No. 2, pp.117–133.

Glueing grids and clouds together

- Chuang, C-C. and Kao, S-J. (2011) 'Adjustable flooding-based discovery with multiple QoSs for cloud services acquisition', *International Journal of Web and Grid Services*, Vol. 7, No. 2, pp.208–224.
- Distributed Systems Technology Centre (2004, March) *E-Research Middleware: The Missing Link in Australia's e-Research Agenda*, Discussion Whitepaper, The Commonwealth of Australia, Department of Education, Science and Training, National Research Infrastructure Taskforce.
- Flahive, A., Taniar, D. and Rahayu, W. (2011) 'Ontology as a Service (OaaS): a case for sub-ontology merging on the cloud', *The Journal of Supercomputing*, pp.1–32.
- Flahive, A., Taniar, D., Rahayu, W. and Apduhan, B.O. (2009) 'Ontology tailoring in the semantic grid', *Computer Standards and Interfaces*, Vol. 31, No. 5, pp.870–885.
- Goodale, T., Jha, S., Kaiser, H., Kielmann, T., Kleijer, P., Merzky, A., Shalf, J. and Smith, C. (2008) *A Simple API for Grid Applications (SAGA)*. Available online at: www.gridforum.org/documents/GFD.90.pdf (accessed on 5 December 2011).
- Hasham, K., Peris, A.D., Anjum, A., Evans, D., Hufnager, D., Huedo, E., Hernandez, J.M., McClatchey, R., Gowdy, S. and Metson, S. (2011) 'CMS workflow execution using intelligent job scheduling and data access strategies', *IEEE Transactions on Nuclear Science (IEEE TNS)*, Vol. 58, No. 3, pp.1221–1232.
- Jeffery, K. and Neidecker-Lutz, B. (2010) *The Future of Cloud Computing: Opportunities for European Cloud Computing beyond 2010*, Expert Group Report Public Version 1.0. Available online at: http://cordis.europa.eu/fp7/ict/ssai/docs/executivesummary-forweb_en.pdf
- Jha, S., Kaiser, H., Merzky, A. and Weidner, O. (2007) 'Grid interoperability at the application level using SAGA', *Proceedings of the 3rd IEEE International Conference on e-Science and Grid Computing (E-SCIENCE '07)*, IEEE Computer Society, Washington, DC, USA, pp.584–591.
- Kertész, A. and Kacsuk, P. (2009) 'Grid interoperability solutions in grid resource management', *IEEE Systems Journal*, Vol. 3, No. 1, pp.131–141.
- Koehler, M., Ruckebauer, M., Janciak, I., Benkner, S., Lischka, H., Gansterer, W.N. (2010) 'A grid services cloud for molecular modelling, workflows', *International Journal of Web and Grid Services*, Vol. 6, No. 2, pp.176–195.
- Kretsis, A., Kokkinos, P. and Varvarigos, E. (2009) 'Developing scheduling policies in gLite middleware', *Proceedings of the 2009th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID '09)*, IEEE Computer Society, Washington, DC, USA, pp.20–27.
- Lee, W., Mcgough, A.S. and Darlington, J. (2005) 'Performance evaluation of the GridSAM job submission and monitoring system', *UK e-Science All Hands Meeting*, pp.915–922.
- Microsoft Windows Azure Platform (2011) *Flexible Open Rock Solid*. Available online at: www.microsoft.com/WindowsAzure (accessed 5 December 2011).
- Nurmi, D., Wolski, R., Grzegorzczak, C., Obertelli, G., Soman, S., Youseff, L. and Zagorodnov, D. (2009) 'The eucalyptus open-source cloud-computing system', *Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID '09)*, IEEE Computer Society, Washington, DC, USA, pp.124–131.
- OCCI (Open Cloud Computing Interface) (2011) *An Open Computing Leading Cloud Standards*. Available online at: <http://www.occi-wg.org/> (accessed 5 December 2011).
- OGF Working Group (2011) *Open Cloud Computing Interface WG (OCCI-WG)*. Available online at: http://www.ogf.org/gf/group_info/view.php?group=occi-wg (accessed on 5 December 2011).
- Open Grid Forum (2011) Available online at: <http://www.gridforum.org/> (accessed on 5 December 2011).
- OpenStack (2011) *Open Source Cloud Computing Software*. Available online at: <http://www.openstack.org/> (accessed on 5 December 2011).
- Rack Space Cloud Hosting (2011) *Open Cloud = Choice*. Available online at: <http://www.rackspace.com> (accessed on 5 December 2011).

A. Anjum et al.

- Redolfi, A., McClatchey, R., Anjum, A., Zijdenbos, A., Manset, D., Barkhof, F., Spenger, C., Legré, Y., Wahlund, L-O., Barattieri, C. and Frisoni, G.B. (2009) 'Grid infrastructures for computational neuroscience: the neuGRID example', *Future Neurology*, Vol. 4, No. 6, pp.703–722.
- SNIA (2011) *CDMI, Common Data Management Interface*, SNIA Technical Position, Cloud Data Management Interface (CDMI) v1.0.1. Available online at: <http://www.snia.org/cdmi> (accessed on 5 December 2011).
- Sun, Q., Wang, S., Zou, H. and Yang, F. (2011) 'QSSA: a QoS-aware service selection approach', *International Journal of Web and Grid Services*, Vol. 7, No. 2, pp.147–169.
- Thain, D., Tannenbaum, T. and Livny, M. (2005) 'Distributed computing in practice: the condor experience', *Concurrency and Computation: Practice and Experience*, Vol. 17, Nos. 2–4, pp.323–356.
- Uehara, M. (2011) 'Composite RAID for rapid prototyping data grid', *International Journal of Web and Grid Services*, Vol. 7, No. 1, pp.58–74.
- van Nieuwpoort, R.V., Kielmann, T. and Bal, H. (2007) 'User-friendly and reliable grid computing based on imperfect middleware', *Proceedings of the ACM/IEEE Conference on Supercomputing (SC'07)*, New York, NY, USA.
- Wikipedia (2011) *OVF, Open Virtualization Format*. Available online at: http://en.wikipedia.org/wiki/Open_Virtualization_Format (accessed on 5 December 2011).