Towards Grid Interoperability

Richard Boardman, **Stephen Crouch**, Hugo Mills, Steven Newhouse, Juri Papay and the OMII-UK team

Open Middleware Infrastructure Institute UK (OMII-UK)
University of Southampton

Abstract

The Grid paradigm promises to provide global access to computing resources, data storage and experimental instruments. It also provides an elegant solution to many resource administration and provisioning problems while offering a platform for collaboration and resource sharing. Although substantial progress has been made towards these goals, nevertheless there is still a lot of work to be done until the Grid can deliver its promises. One of the central issues is the development of standards and Grid interoperability. Job execution is one of the key capabilities in all Grid environments. This is a well understood, mature area with standards and implementations. This paper describes some proof of concept experiments demonstrating the interoperability between various Grid environments.

1. Introduction

Over the last decade the Grid computing paradigm has undergone several changes that reflect the evolution of this concept and also the changing perceptions of the Grid. Originally, the term Grid that was coined in the mid-1990's described the vision of a shared computing infrastructure for researchers [Foster1]. This vision was later refined and described the Grid as an infrastructure for "coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations" [Foster2].

Alongside the evolution of the technology, its use within different communities has also changed. The Grid is no longer seen just as a mechanism for federating supercomputers. Different user communities see the Grid from a different perspective. For the particle physics community, the Grid is a massive, loosely coupled, distributed computing environment with extensive computing capability, communication bandwidth and data storage capacity. For bio-informaticians the Grid is a global virtual federated database of experimental data, research papers and laboratory records. The bio-informaticians would also like to see this database annotated with metadata and comments that enrich the information content. Engineers, and other disciplines, see the Grid as a powerful tool to supporting the exploration of "what if" experiments through parameter sweeps across their computer models.

Despite the different perceptions, perhaps all communities agree that the Grid provides secure virtualisation of resources (i.e. computing power, communication bandwidth, data, information and knowledge), and enables collaboration and establishment of virtual organisations. This perception provides the "smallest common denominator" on which the community can agree. Naturally different communities will add extra layers to this "common denominator" as a reflection of their needs and interests.

Since the idea of Grid computing has gained a foothold both in academia and industry a multitude of Grid environments have been developed. These environments work reasonably well within their own boundaries; however there is little interoperability between them. The main reasons are the diversity of software platforms, the various policies governing the access and use of distributed resources and the apparent lack of adherence to the emerging standards.

Over the recent years substantial effort has been invested in the development of Grid standards. There are several internationally recognised bodies that are active in this area that provide significant contributions, in particular the Open Grid Forum (OGF) [OGF], the Organisation for the Advancement of Structured Information Standards (OASIS) [OASIS], the World Wide Web Consortium (W3C) [W3C] and the Distributed Management Task Force (DMTF) [DMTF] as well as others.

This illustrates the increasing focus of the Grid within the web services arena, and the importance of standards. Such developments in this area indicate that web-service based technology is likely to become the foundation of Grid middleware. One of the issues is the large number of various proposals of standards and the extensive overlap and even frequent inconsistency between the draft documents. This large number of proposals indicates the intensive development, however on the other side the fluidity of standards significantly hinders the implementation effort that ultimately may slow down the adoption of Grid technologies by a wider user community.

There is no doubt that interoperability of Grids can offer huge benefits for the user community. By this community we understand different categories such as e-Infrastructure providers, e-Science users and e-Science application developers. For the e-Infrastructure providers, interoperability would enable easier deployment and management of software distributions. For the e-Science users it would give the freedom to choose services that are deployed in different Grids. The selection of services would be based on their functionality rather than on their deployment in a particular Grid. Interoperability for the e-Science application developers would provide portability of their applications across multiple Grids rather than limiting the scope of applications for a specific Grid.

Interoperability assumes the use of interfaces described by generally accepted standards. Without compliance with the standards we cannot talk about interoperability. This idea has been strongly supported and systematically implemented by the Grid Interoperability Now (GIN) initiative that is a community activity within the Open Grid Forum (OGF) [OGF]. The idea is to establish so-called interoperability islands that would cover the key services such as Job Submission, Information Services, Storage Management, Accounting, Job Monitoring, Database Access and Virtual Organisation Management. The cornerstone of interoperability for Job Submission is adherence to the WS-GRAM [WSGRAM], Job Submission Description Language (JSDL) [JSDL] and Basic Execution Service (BES) [BES] standards. The initial target for Federated Information Services is to achieve full interoperability between popular Grid environments such as EGEE [EGEE], OSG [OSG] and NAREGI [NAREGI]. The effort for Storage Management standardisation focuses on the interoperability of the SRB [SRB] and SRM [SRM] products. A generally accepted standard for Data Movement is GridFTP [GridFTP] and for Authorisation the VOMS [VOMS] specification.

There is a significant effort invested into standardisation of building and testing infrastructures. The main candidate for the building infrastructure is the Metronome [Metronome] suite and for testing the ETICS [ETICS] environment.

2. Build and Test Configuration Management

Grids represent complex distributed systems with diverse computing platforms, software and experimental equipment. Easy and transparent building and deployment of services in such environments is one of the key requirements. In order to meet this requirement it is essential to have an infrastructure that complies with a generally accepted standard and provides a cross-platform environment for building and testing. Additional requirement to this infrastructure is to provide support for testing for conformance and interoperability of different grids.

ETICS is a CERN-led effort that aims at leveraging the capabilities of the Metronome framework (formerly the NMI Build and Test framework) across Condor [Condor] to control the management of software builds and testing procedures. In principle, it will allow developers to create configurations for projects and keep a

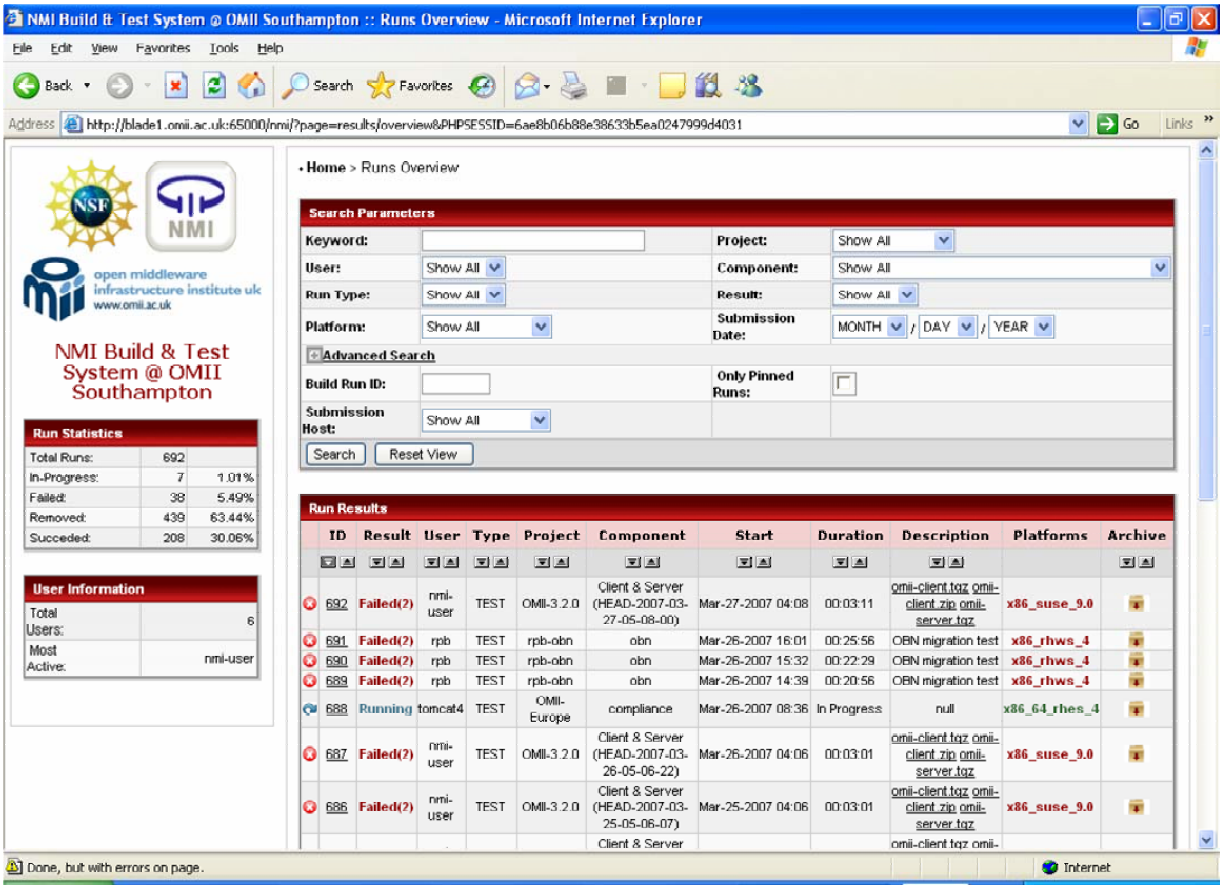

Figure 1 - Metronome building and testing environment

historical account of build and test runs, with the capacity for jUnit [jUnit] reporting presentation. In the wider scope the ultimate aim is to develop a single access point represented by a portal that would enable the component developers to submit their software and automatically build and test it. The Southampton team at OMII-UK recently developed a software that shows the capabilities of the ETICS environment. This software includes the submission of tasks to a resource manager (e.g. GridSAM [GridSAM]) within a testing suite (Metronome/ETICS) designed to measure the level of compliance with those standards. Figure 1 represents a snapshot from the Metronome environment.

With the deployment of the ETICS web application at OMII-UK in Southampton, it is now possible to develop and submit software builds and run tests orchestrated from the central ETICS database at CERN. Code metrics are now calculated for each build and test, including code complexity, style checking, test coverage and SLOC counts. These reports are presented as formatted HTML, and are also provided in a standard XML format. Automatic package production is also provided through the ETICS system. Eventually it will be possible to submit jobs through the Southampton ETICS web application and run them across the wide variety of platforms available at CERN and Wisconsin's grandcentral machine pool.

3. Emerging job submission standards

There are two key elements of Job Submissions for Grids. One is the interface and the other is the Job Submission Description Language (JSDL). The interface for creating and managing a job is described by the Open Grid Services Architecture (OGSA) Basic Execution Service (BES) standard. This interface provides a single access point for submitting jobs. BES is not intended to handle anything more complex

BES-Management Port-type	
StopAcceptingNewActivities	Request that the BES stop accepting new activities
StartAcceptingNewActivities	Request that the BES start accepting new activities
BES-Factory Port-type	
CreateActivity	Request the creation of a new activity
GetActivityStatuses	Request the status of a set of activities
TerminateActivities	Request that a set of activities be terminated
GetActivityDocuments	Request the JSDL documents for a set of activities
GetFactoryAttributesDocument	Request XML document containing BES properties
BES-Activity Port-type (optional)	
GetStatus	Request the status of an activity
Terminate	Request that an activity be terminated
GetDocument	Request the JSDL document for an activity
GetActivityAttributesDocument	Request XML document containing activity properties

Table 1 – BES port types and functions

than simple job submission – tasks such as scheduling or resource discovery are left to other services, which should be used to identify the resource on which to execute. The interface represents a collection of methods that are exposed to the client application. Behind this interface there is the execution engine, running on a cluster or a network of computers or representing shared experimental equipment. BES is currently in its public comment period, and will shortly be published as an OGF standard.

BES defines three Web Service port types: for Management, Activity Execution and an optional Activity Management port type. Each port type contains a limited number of simple methods (see Table 1) which are nevertheless sufficient to cover the lifecycle of a job. BES also defines a simple state model for jobs, and methods for extending the state model and other aspects of the system so as to be able to handle additional operational semantics in an implementation-specific manner.

There is a simplified form of BES, representing its core functionality, embodied in the High Performance Computing Profile (HPC-Profile) [HPCP] standard. HPC-Profile removes many of the complexities needed by a full BES implementation, such as the need for WS-Security support. The security framework for HPC-Profile instead uses HTTPS as a communications channel, and certificates for user authentication. There are various implementations of the HPC-Profile interface, from most of the major developers and vendors of Grid software, including OMII-UK [OMII-UK], Globus [Globus] and Unicore [Unicore]. The OMII-UK implementation is based on the new GridSAM 2.0.1 release, the first to support the BES port types.

Over the years JSDL has become a *de-facto* standard for the description of computational jobs. One of the key elements of JSDL is the use of XML syntax for specifying the job parameters and details of execution. Describing jobs in XML provides portability in other words job description documents can be used across different Grids.

The OMII-UK Southampton team has implemented several scenarios for representing the key features of Job Submission Service and its interoperability with different Grid environments. The scenarios use a variable number of clients to access multiple service providers. The work shows interoperability between resource managers and web service platforms using the HPC-Profile. The HPC-Profile is a proposed standard for Grid interoperability in HPC environments that references existing specifications including: JSDL, BES and WS-I Basic Profile. The OGSA HPC-Profile allows application and middleware software providers to target multiple HPC

Client/HPCP implementation	UVa .NET	Microsoft	Platform	Globus	Unicore	Genesis II	GridSAM	CROWN	WMProxy	CREAM
UVA .NET	—	*done*	*done*			*done*				
Microsoft HPC Group	*done*	—	*done*	*done*			*done*	*done*		
Platform	*done*	*done*	—	*done*		*done*	*done*	*done*		
Globus	*done*	*done*	*done*	—			*done*	*done*		*done*
Unicore					—					
Genesis II	*done*	*done*	*done*	*done*		—	*done*	*done*		
GridSAM	*done*		*done*	*done*		*done*	—	*done*	*done*	*done*
CROWN	*done*	*done*	*done*	*done*		*done*	*done*	—	*done*	*done*
gLite WMProxy									—	
gLite CREAM								*done*	*done*	—

Table 2 – Interoperability between clients and different HPC-Profile implementations

systems via a single protocol. For the compliance testing the following sequence has been used:

- get factory attributes document
- create job
- query the status of that job
- show job output
- get job's JSDL document
- terminate job

This sequence covers the core functionality of HPC-Profile and checks whether the client can communicate with different implementations of the HPC-Profile interface. Each method from the above mentioned sequence has been individually tested on different implementations of the HPC-Profile interface. For brevity, in this paper we provide only a summary table for the "*show job output method*" that summarises the output of interoperability experiments presented at Supercomputing 2006 [Super06].

The OMII-UK Southampton team has been involved in the development of GridSAM and CROWN [CROWN] implementations of HPC-Profile interface. From Table 2 we can see that the CROWN client is interoperable with all HPC-Profile implementations and the GridSAM client is interoperable with all except the Microsoft implementation of HPC-Profile interface.

4. Full work cycle

One of the benefits of interoperability is the freedom to select services based on the required functionality rather than on their availability in a specific Grid environment. This demonstration shows the full work cycle that consists of service discovery, workflow composition, launching the workflow and collecting the results. In this case we used the Grimoires (Grid Registry with Metadata Oriented Interface: Robustness, Efficiency, Security) [Grimoires] registry for service discovery and selection, Taverna [Taverna] for worflow design and GridSAM [GridSAM] for job submission and PlotWS for visualisation.

The left panel of Figure 2 provides the graphical representation of the workflow. The middle panel is the textual workflow editor. The bottom of the right panel gives a hierarchical list of available services. The top panel provides input/output and process monitoring functionalities.

Grimoires enables storage of service descriptions, distributed queries, WSDL documents and workflows. This registry also provides facilities for semantic annotation of information. Grimoires is fully UDDIv2 [UDDIv2] standard compliant. In addition to the UDDIv2 interface, Grimoires also provides some other interfaces, such as a metadata interface and a WSDL interface, which allow clients to publish and inquire over metadata and WSDL-related data, respectively. All the data published through various interfaces are internally represented as RDF triples, which can be queried and reasoned about in a uniform way.

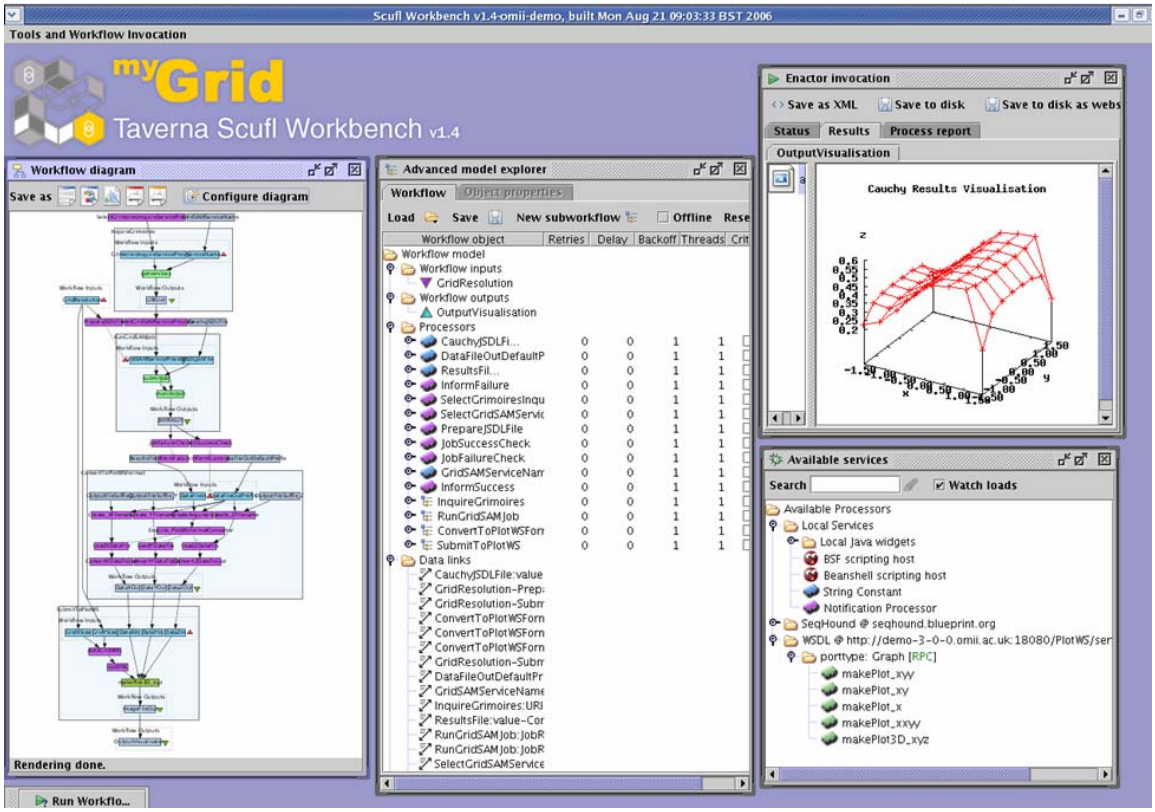

Figure 2 – Interaction between individual services Taverna, Grimoires, GridSAM and PlotWS

The Taverna [Taverna] tool aims to provide a language and software tools to facilitate easy use of workflow and distributed computer technology within the e-Science community. GridSAM is an open-source job submission and monitoring service. GridSAM installs on top of the WS-Security (authentication) layer provided by the OMII-UK WS container and enables users to execute jobs on the OMII-UK server that may have a variety of data input and output requirements. GridSAM implements JSDL and HPC-Profile standards.

In this scenario Taverna's role is to coordinate the workflow between the GridSAM, Grimoires and PlotWS secure web services. The organisational and technology-interface mechanisms provided by Taverna are used to achieve a set of scientific and operational objectives, and thus illustrates, through Taverna, how such technologies can be made to easily work together.

5. Job Brokering using the CROWNGrid Scheduler

As part of the OMII-UK's work evaluating and integrating technology for OMII-Europe [OMII-Europe] we have shown interoperability between the CROWN and OMII-UK software environments for job brokering, resource discovery and job execution. In this demonstration, the CROWN [CROWN] scheduler, deployed within the OMII-UK web services container, accepts JSDL job descriptions, which then requests any usable job services from a Grimoires registry. It is then able to select an execution service based on its scheduling policy and submit the job to instances of either GridSAM or the CROWN job services for execution (Figure 3).

Again, we see the benefits of Grid vendors converging on a common set of standards. This integration was possible thanks to the fact that both the OMII-UK and CROWN platforms use the same set of core interoperable standards such as JSDL.

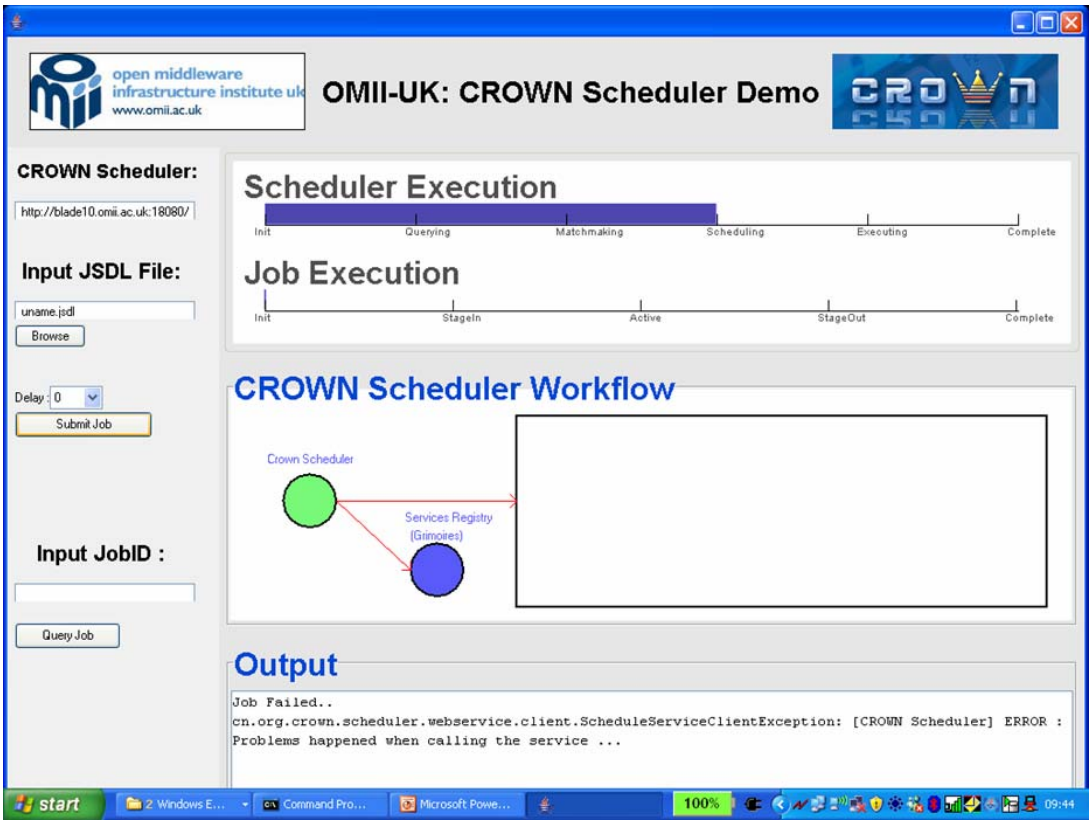

Figure 3 - Job Brokering using the CROWNGrid Scheduler

6. Summary

The main points detailed in this paper can be summarised as follows:

- the benefits of standardisation from the point of view of different user communities
- the significance of the Metronome platform as a building and testing environment
- the current status of Job Submission standards
- a summary of the endpoint compatibility verification experiments of different HPC-Profile implementations
- the importance of interoperability in respect to workflow design, service selection and job execution in the Grid environment
- the use of the CROWN scheduler for submitting jobs into CROWNGrid and OMII-UK Grid environments

There is no doubt that in case of the Grid we are dealing with a new phenomenon of unprecedented complexity that requires the solution, not only of technical problems, but of organisational and even political issues as well. Following on from the first wave of projects, which exploited new networking infrastructures, experimental test-beds, middleware and application software, we have a far greater understanding of the key issues. The next phase of activity will concentrate on interoperability and running applications that clearly demonstrate the benefits of the Grid. There is a strong commitment to allocate more resources for this purpose that will certainly get us closer to the materialisation of the Grid promises.

At present experimentation is being conducted with a wide range of Grid middleware products. Many of these products are still at the prototype stage and only a few of them have approached industrial strength quality. In the future it is likely that multiple Grid middleware products will be used and therefore the problem of interoperability is of key importance. It is also expected that the Grid middleware will be based on open standards, open source and royalty free schemes.

7. References

[Super06]: Info about interoperability demos: https://forge.gridforum.org/sf/wiki/do/viewPage/projects.ogsa-hpcp-wg/wiki/HomePage, November 16 2006

[Foster1]: Foster, I., Kesselman, C., (eds), "The Grid: Blueprint for a New Computing Infrastructure", Morgan Kaufmann, 1999.

[Foster2]: Foster, I., Kesselman, C., Tuecke, S., "The Anatomy of the Grid: Enabling Scalable Virtual Organisations,International Journal of High Performance", Computing Applications 15 (3), 200-222, 2001.

[OASIS]: Organisation for the Advancement of Structured Information Standards (OASIS), http://www.oasis-open.org

[DMTF]: Distributed Management Task Force (DMTF), http://www.dmtf.org

[Grimoires]: http://www.ecs.soton.ac.uk/research/projects/grimoires

[UDDIv2]: http://www.oasis-open.org/committees/uddi-spec/doc/tcspecs.htm

[GridSAM]: http://gridsam.sourceforge.net/2.0.0-SNAPSHOT/index.html

[Taverna]: http://taverna.sourceforge.net/

[OGF]: Open Grid Forum, http://www.ogf.org

[WSGRAM]: http://www.globus.org/toolkit/docs/4.0/execution/wsgram/

[JSDL]: www.gridforum.org/documents/GFD.56.pdf

[BES]: http://www.ogf.org/pipermail/ogsa-bes-wg/attachments/20060906/c1849ef3/attachment-0003.doc

[EGEE]: http:// www.eu-egee.org/

[OSG]: http://www.opensciencegrid.org/

[NAREGI]: http://www.naregi.org/

[SRB]: http://www.sdsc.edu/srb/index.php/Main_Page

[SRM]: http://sdm.lbl.gov/indexproj.php?ProjectID=SRM

[GridFTP]: http://dev.globus.org/wiki/GridFTP

[VOMS]: EU DataGrid, VOMS Architecture v1.1, http://grid-auth.infn.it/docs/VOMSv1_1.pdf

[ETICS]: http://etics.web.cern.ch/etics/

[jUnit]: http://www.junit.org

[Condor]: http://www.cs.wisc.edu/condor/

[HPCP]: http://www.ogf.org/Public_Comment_Docs/Documents/Feb-2007/HPC_Basic_Profile_v1.0.pdf

[Metronome]: http://nmi.cs.wisc.edu/node/90

[OMII-UK]: http://www.omii.ac.uk

[OMII-Europe]: http://www.omii-europe.org

[Globus]: http://www.globus.org

[Unicore]: http://www.unicore.org/

[WS-I]: http://www.ws-i.org/

[Crown]: http://www.crown.org.cn/en/