



**omii europe**

open middleware infrastructure institute



EU project: RI031844-OMII-Europe

Project no: **RI031844-OMII-Europe**

Project acronym: **OMII-Europe**

Project title: **Open Middleware Infrastructure Institute for Europe**

Instrument: **Integrated Infrastructure Initiative**

Thematic Priority: **Communication network development**

## **D:JRA1.7 1<sup>st</sup> Yearly Report on the Accounting Task**

Due date of deliverable: 30 April 2007

Actual submission date: 17 May 2007

Start date of project: **1 May 2006**

Duration: **2 years**

Kungliga Tekniska Högskolan (KTH)

Revision: 0.7

<b>Project co-funded by the European Commission within the Sixth Framework Programme (2002-2006)</b>		
<b>Dissemination Level:</b>		
<b>PU</b>	Public	X
<b>PP</b>	Restricted to other programme participants (including the Commission Services)	
<b>RE</b>	Restricted to a group specified by the consortium (including the Commission Services)	
<b>CO</b>	Confidential, only for members of the consortium (including the Commission Services)	

## Document Control Sheet

<b>Document</b>	<b>Title:</b> 1 <sup>st</sup> Yearly Report on the Accounting Task	
	<b>ID:</b> D:JRA1.7	
	<b>Version:</b> 1.0	<b>Status:</b> Final
	<b>Available at:</b> TBD	
	<b>Software Tool:</b> Microsoft Word 2004	
	<b>File(s):</b> DJRA1.7.doc, DJRA1.7.pdf	
<b>Authorship</b>	<b>Written by:</b>	Gilbert Netzer (KTH)
	<b>Contributors:</b>	Michele Pace (INFN) Morris Riedel (FZJ) Fredrik Hedman (KTH)
	<b>Reviewed by:</b>	Moreno Marzolla (INFN) Mike Mineter (UEDIN)
	<b>Approved by:</b>	Technical Committee

## Document Status Sheet

Version	Date	Status	Comments
0.1	27 April 2007	Draft	All sections except UNICORE written. (GN)
0.2	2 May 2007	Draft	Compacted planning section, added overview figure. (GN)
0.3	3 May 2007	Draft	Reformatted 2nd year plan into table. (GN)
0.4	13 May 2007	Draft	UNICORE sections written, small corrections (MR)
0.5	17 May 2007	Done	Added Partners and Efforts Section. (FH)
0.6	21 May 2007	Done	Formatted document, small corrections. (GN)
0.7	31 May 2007	Final	Edited to meet comments by reviewers. (GN)

# 1 Executive Summary

The document contains the description of the work done during the first 12 month by the accounting task of the JRA1 re-engineering of services activity of the OMII-Europe project. 7 partners (KTG, INFN, FZJ, FLE, BU and CNIC) contribute to the task. The task is coordinated by KTH. A half time report of the whole task is given and a plan for the second year is presented. The report presents a summary of the accomplished milestones, more detailed descriptions of the achievements can be found in the milestone documents M:JRA1.4[1] and M:JRA1.5[2] written by the activity. The report is mostly intended as an overview for management and reviewers.

The overall goal of the accounting task is to augment the accounting systems of the three Grid middleware stacks gLite, Globus toolkit and UNICORE with a Resource Usage Service (RUS) compliant standard interface to allow the exchange of accounting information.

During the first year of the task, an analysis of the existing systems was undertaken and a preliminary design for realizing the task goal was created and published as milestone M:JRA1.4[1] in September 2007 (M5). The task then progressed with the implementation of RUS interfaces for the Globus and UNICORE stack according to original plans. For the gLite stack, the task plan was altered to adjust to staff changes in coordination with project management and the commission. The work planned on the information providers (usage record generators) for gLite during year 2 was conducted in the first year to take advantage of existing prototypes and knowledge while the work on the interface was postponed to year 2. This also allowed the activity to demonstrate interoperability between gLite and Globus components at the M12 project meeting. The alpha release milestone M:JRA1.5[2] was completed in time during April 2007 (M12). The RUS interface for UNICORE was successfully demonstrated with the LLView resource usage monitoring application at the M12 all hands meeting in Jülich. This work was published as a paper and presented at the German e-Science Conference 2007.

The first milestone (M:JRA1.4) was completed with a slight delay (M5 instead of M3) due to ramp-up effects during the task start-up. This was also brought to the attention of the commission and the new delivery date September 2006 was set and met by the task. The alpha release milestone goal had to be slightly altered by redistribution work between first and second project year, with the benefit of demonstrating early interoperability. The modified M:JRA1.5 milestone was completed on time in April 2007 (M12). Overall task progress is according to the project plan.

# Table of Contents

1	Executive Summary .....	3
2	Activity Overview .....	5
3	Partners and Efforts .....	6
4	Preliminary Design Document (M:JRA1.4) .....	7
5	Feature Complete Alpha Implementations (M:JRA1.5) .....	7
	Changes in the Milestone Objective .....	7
	RUS Specification Draft .....	7
	DGAS .....	8
	SGAS .....	8
	UNICORE .....	8
	LLView use case application using UNICORE RUS interface .....	9
	Early Interoperability Preview .....	9
6	Community Involvement .....	10
	Usage Record Working Group (UR-WG) .....	10
	Resource Usage Service Working Group (RUS-WG) .....	10
7	Outlook for Project Year 2 .....	10
	Planned Interactions .....	11

## 2 Activity Overview

Accounting systems record information about the usage of various resources like computational jobs, disk storage or network usage. The collected information is then digested into summary reports, bills or usage statistics showing historic usage consumption. This gives stakeholders like users, resource owners and managers the necessary information to assess and monitor the current status or the fulfilment of contracts or agreements. In addition many accounting systems have options to prevent over-usage of resources by single users (quota enforcement) so that agreements with other users can be fulfilled. In the Grid context this becomes even more important in the light of the growing complexity and number of agreements between resource owners and consumers.

With projects spanning multiple Grid deployments, exchange of accounting information becomes a problem. For instance the Nordic Tier-1 of the LHC Computing Grid (LCG) is using resources from the national grids of the Scandinavian countries. In this case the accounting information has to be reported back to the central Nordic DataGrid Facility that is responsible for the correct operation of the Tier-1. Another use-case in the OMII-Europe context is the exchange of usage monitoring information within the German national Grid (D-Grid). OMII-Europe tries to address that problem by modifying existing accounting systems to follow standardized interfaces and data formats developed by the Open Grid Forum (OGF) community process which should reduce the need for incompatible ad-hoc solutions.

The approved Description of Work [3] defines the following objectives for the accounting task in JRA1:

The GGF has recently published an initial draft specification for exchanging information on resource usage (RUS). OMII-Europe will unify accounting information on EGEE, UNICORE and GLOBUS by integrating and supporting a common standardised specification such as RUS on these Grid distributions. In year 1 existing accounting information systems (DGAS/SGAS/UNICORE) will be augmented with the RUS (or equivalent) interface. In year 2, the Grid software distributions will be instrumented with the necessary information providers, that is, the clients register and report accounting information back to the RUS. In UNICORE, this latter functionality already exists.

Table 1 summarizes the activity milestones and deliverables. The activity deliverables will also be included into the JRA1 deliverables D:JRA1.1 and D:JRA1.2.

Figure 1 gives an overview of the Resource Usage Service (RUS) related accounting system components in the three middleware stacks. All three systems collect usage information via metering components that are deployed at the resources. This information is then passed on to the service level, where it is stored and possibly digested. Users then use the provided client applications to query the services for the stored data. If quota enforcement is in place, the services

No.	Title	Date	Completed
M:JRA1.4	Preliminary design document	M5	M5
M:JRA1.5	Feature complete alpha release	M12	M12
M:JRA1.6	First release of integrated components	M20	
D:JRA1.7	First yearly report on accounting task	M12	M12
D:JRA1.8	Second yearly report on accounting task	M23	

*Table 1: Overview of the Task Milestones and Deliverables*

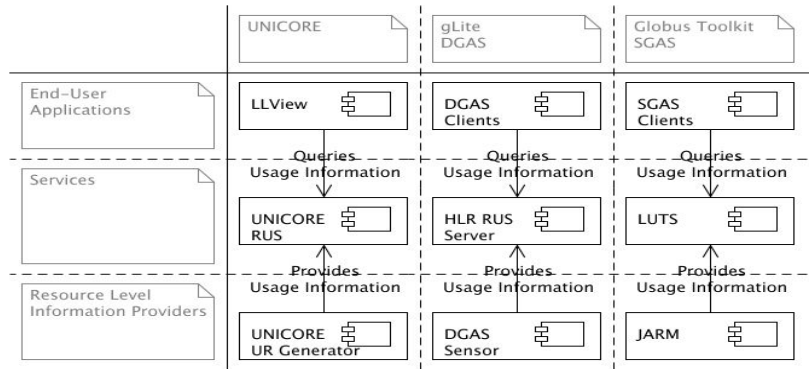


Figure 1: Overview over the JRA1 accounting components in the three middleware stacks UNICORE, gLite and Globus (columns). Each stack has components in three different layers (end-user applications, services and information providers).

also keep track of the remaining quotas and report back to the metering components if the quotas are exceeded. The metering components then give a recommendation back to the local batch system at the resource if the job should be executed or not.

### 3 Partners and Efforts

This work involves 10.5 staff years, KTH 2.0 (2.0 funded), INFN 3.3 (1.9 funded, 1.4 unfunded), FZJ 0.2 (0.2 funded), FLE 0.6 (0.6 funded), UEDIN 2.4 (0.4 funded, 2.0 unfunded), BU 1.0 (1.0 unfunded) and CNIC 1.0 (1.0 unfunded). KTH and INFN bring in RUS expertise. FZJ and FLE bring in UNICORE expertise, UEDIN brings in Globus expertise. BU and CNIC will port the accounting systems to CROWNGrid. KTH will lead this task.

During the start-up phase of the project, the task was temporarily coordinated by INFN until KTH, the designated coordinating partner, was able to take over this responsibility after initial staff shortage. The handover took place during the month 6 face-to-face meeting in Southampton. KTH has been coordinating the activity since then. Table 2 shows the efforts from participating partners during the first project year in staff months and the fraction of effort that should have been spent

Org.	KTH		INFN		FZJ		FLE		UEDIN		BU		CNIC	
Efforts (F/U)	1,00	0,00	0,95	0,70	0,10	0,00	0,30	0,00	0,20	1,00	0,00	0,50	0,00	0,50
May-06	0,00		1,00	0,00	0,10		0,00		0,10	0,00		0,00		0,00
Jun-06	0,00		1,00	0,00	0,10		0,00		0,00	0,00		0,00		0,30
Jul-06	0,00		1,00	0,70	0,10		0,20		0,10	0,00		0,00		0,00
Aug-06	0,00		1,00	0,70	0,10		0,20		0,20	1,00		0,00		0,00
Sep-06	0,00		1,00	0,70	0,10		0,20		0,20	1,00		0,00		0,00
Oct-06	1,20		1,00	0,70	0,10		0,20		0,20	1,00		0,10		0,10
Nov-06	1,00		0,30	0,50	0,10		0,20		0,20	1,00		0,10		0,10
Dec-06	1,00		0,00	0,50	0,10		0,00		0,20	1,00		0,00		0,10
Jan-07	1,35		0,00	0,50	0,10		0,20		0,20	1,00		0,00		0,00
Feb-07	1,35		0,66	0,50	0,10		0,20		0,20	1,00		0,00		0,15
Mar-07	1,95		1,00	0,50	0,10		0,20		0,20	1,00		0,00		0,10
Apr-07	1,95		1,00	0,50	0,10		0,20		0,20	1,00		0,00		0,00
<b>Sum (Months)</b>	<b>9,80</b>		<b>8,96</b>	<b>5,80</b>	<b>1,20</b>		<b>1,80</b>		<b>2,00</b>	<b>9,00</b>		<b>0,20</b>		<b>0,85</b>
<b>Fraction</b>	<b>0,82</b>		<b>0,79</b>	<b>0,69</b>	<b>1,00</b>		<b>0,50</b>		<b>0,83</b>	<b>0,75</b>		<b>0,03</b>		<b>0,14</b>

Table 2: Partner efforts. Overview of efforts during the first project year, in staff months. (F/U) is funded and unfunded respectively. The fraction represents the fraction of the effort for the first 12 months that has actually been spent.

during the first year. Due to a combination of initial staff shortage and employees leaving the project it is clear that the effective staff efforts will need to increase during the coming 12 months. Furthermore, the extent of actual participation from the BU and CNIC is under investigation. However, despite these staff fluctuations most of the planned tasks have been completed for the first year. The main remaining effect is that the order of some of the tasks has been changed.

#### **4 Preliminary Design Document (M:JRA1.4)**

The first milestone of this task was to present a plan on how the three Grid middleware stacks gLite, Globus and UNICORE should be augmented with a RUS interface to allow for the exchange of accounting data.

The original delivery date of the 31st of July 2006 (M3) was unfortunately unrealistic due to ramp up effects in the beginning of the project. Therefore the deadline was in cooperation with the commission revised to the 30th September 2006 (M5). Despite that delay, it was anticipated that the delivery date of the next milestone would not be affected. The milestone document [1] was completed and delivered to the European Commission on the 30th of September 2006.

An analysis of the current architecture of the accounting systems of the three middlewares and of the status of the emerging RUS specification was successfully conducted. Also a principal design for the integration of RUS functionality in the three accounting components was presented. The solutions were tailored to the reuse of pre-existing components in the middlewares to allow seamless integration of the proposed RUS implementations.

#### **5 Feature Complete Alpha Implementations (M:JRA1.5)**

The goal of the second milestone (M:JRA1.5) is to provide alpha implementations of the RUS interfaces for the middleware stacks based on the M:JRA1.4 milestone [1]. A detailed technical description of the modifications can be found in the M:JRA1.5 milestone document [2].

##### **Changes in the Milestone Objective**

During this milestone, the project partners had to adapt to a key staff change at INFN, the partner responsible for developing the server component for the gLite middleware. To prevent any delay, the activity partners decided in close collaboration to adjust the task planning to the new circumstances.

The development of the server components for DGAS was postponed until year 2 and efforts were focused on the implementation of a RUS compliant sensor that would convert usage information from the local batch system to UsageRecord format and deliver this information to one of the RUS server components developed by the other two project partners. This work was originally planned for the second year. The solution was favourable because a UsageRecord generator prototype was already under development at INFN, which could be used as basis for the sensor implementation.

As an additional benefit, this new plan provides early evidence of the interoperability of the RUS components of different middleware stacks since the output of the DGAS sensor will be used to feed data into the SGAS and UNICORE server components.

##### **RUS Specification Draft**

Since the RUS specification [4] is not yet finalized, the project decided to use the available draft version 17, which has previously undergone a public comment phase. During the further development of the project, the team will assess if a newer draft version of the specification should be used when it becomes available.

## **DGAS**

During the writing of the M:JRA1.4 milestone, a more detailed design plan [5] was created for DGAS as a basis for the server implementation.

As defined in the adapted plan for the M12 milestone, development in DGAS focused on the UR generator and a client to upload the generated UsageRecords into a RUS server. This facilitates the population of a RUS service with usage information from the computing resources. Additionally a client to test the other RUS functionality (query/modify/delete) was written.

During the development of the RUS client a very simple RUS server has been developed mainly to test early design decisions for the DGAS RUS server. The purpose was to test gSoap as a Web service development tool in respect to RUS service descriptors; to test the DGAS client during the development and to make cross interoperability tests between SGAS and DGAS implementations. Currently only one method (the insert) of the service is implemented.

## **SGAS**

Before implementation started for SGAS a more detailed plan [6] was written and published on the project Wiki. This document will be updated during the second year when the implementation stabilizes. The plan proposes to add a RUS interface to the LUTS component of the SGAS. This component is responsible for the storage of UsageRecords.

During the first phase of the augmentation of the LUTS component, most effort was consumed by the changes to the existing database structure and legacy interface to create the audit information required by the RUS interface.

In the second phase of implementation, a RUS interface was added to the LUTS component. The resulting alpha version is a very simple but complete RUS server implementation. Command line test clients have also been written for all RUS methods to be able to test server functionality. However some features are not yet completely implemented or missing. In particular error handling and security is only very rudimentary.

## **UNICORE**

The development of the alpha version of the RUS service within UNICORE 6 followed the initial plan of the milestone M:JRA1.4 [1]. In particular, the service schemas of the draft 17 specification were used to create the necessary RUS interface bindings within UNICORE 6. The RUS service is deployed within the WS-RF compliant hosting environment (based on Jetty & xFire SOAP engine) of UNICORE 6 and thus benefits from the strong security architecture of UNICORE 6. In more detail, since the service is deployed in the same service container as the OGSA-BES and WS-DAI interfaces, the RUS service benefits from the already achieved alpha prototype VOMS integration into UNICORE (JRA1 VOM). Hence, the RUS service within in UNICORE provides automatically authentication (via the UNICORE Gateway) and prototype VOMS authorization (based on SAML & XACML).

In parallel, the overall UNICORE 6 evolved to a beta implementation that will be used in production in June 2007, therefore the next month will focus on the maturing of the UNICORE RUS service to provide this service in production quality as soon as possible for end-users out of OMII-Europe, for instance within DEISA and D-Grid. In the context of JRA3 – Task 2 Infrastructure Integration, the RUS interface will participate, along with DGAS and SGAS RUS interfaces in interoperability scenarios. The German national Grid D-Grid project D-Mon (monitoring) project has already stated interest in using the RUS interface of UNICORE 6 for resource level monitoring.

## LLView use case application using UNICORE RUS interface

On the M12 meeting the LLView resource usage monitoring application was used to demonstrate the availability of the RUS interface within UNICORE. In particular, a Web services-based PERL RUS client within the LLView application was used to extract the usage records from the RUS interface of UNICORE. In this context, the RUS interface was used to extract the up-to-date usage information to allow a usage monitoring of the JUMP supercomputer at FZJ. This work was successfully published and demonstrated at the German e-Science Conference 2007 in Baden-Baden Germany. In the context of interoperability scenarios of JRA3 – Task 2 Infrastructure Integration, it is planned to demonstrate the usage of LLView on top of SGAS and DGAS in the second year as well.

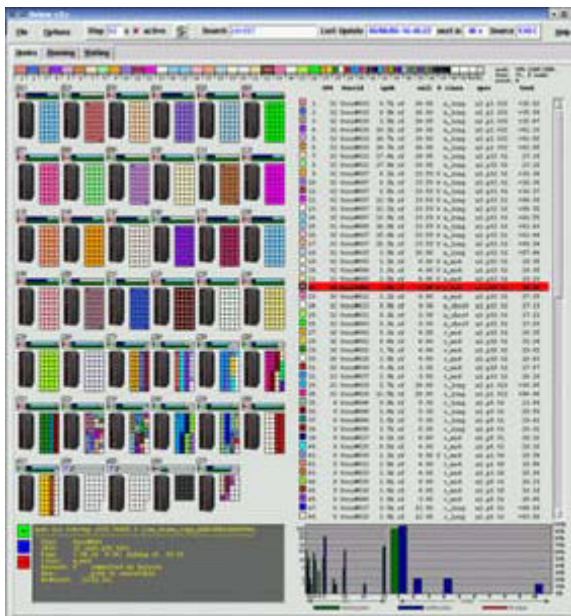


Figure 2: LLView application gets the up-to-date usage information from the RUS interface alpha implementation of UNICORE 6. A usage on top of SGAS and DGAS is planned.

## Early Interoperability Preview

On the M12 meeting and at OGF20, the task demonstrated early interoperability between the alpha versions of the DGAS and SGAS components. As shown in figure 3 usage records were generated

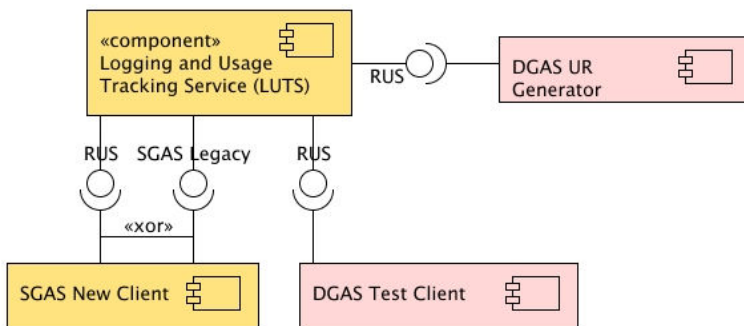


Figure 3: Set-up for the interoperability between DGAS and SGAS alpha releases. SGAS components are shown in yellow, DGAS components in pink. Provided interfaces are shown using balls, used interfaces are shown with open sockets. The LUTS can be queried by the SGAS clients either using the legacy interface or the RUS interface. The DGAS client can only use the RUS interface.

from canned LSF log entries and inserted by the DGAS RUS client into the SGAS LUTS server. The records could then be queried using either the DGAS client or the SGAS test client with both the legacy and the RUS interface.

## **6 Community Involvement**

Since the RUS and UR specifications are continuing to evolve, the project decided to participate in the Open Grid Forum community in order to provide feedback from the implementation work and contribute to the development of the RUS and UR specifications.

### **Usage Record Working Group (UR-WG)**

During the 18th Global Grid Forum, held in Washington DC in September 2006, Rosario M. Piro proposed changes to the Usage Record Format [7]. The proposals were made in cooperation with the EGEE project. Recently, the usage of LLview on top of UNICORE 6 RUS interfaces reveals that the UR specification 1.0 is not fine-grained enough for very precise accounting and resource usage monitoring. Therefore, the fine-grained additions made within this activity to support LLview will be contributed to the development of the UR specification 2.0 within OGF.

### **Resource Usage Service Working Group (RUS-WG)**

Before the 19th Open Grid Forum in Chapel Hill in January 2007, two members of the OMII-Europe project were appointed chairs of the RUS working group. Also an update to the working group charter was presented to the area directors.

During OGF19, a RUS working group session was held to plan further steps in the specification process. Plans were proposed for reformulating the RUS specification in a more technology independent way so that it would become easier to adapt to new transport technologies such as WS-Management. Also plans were presented for writing an advanced RUS specification to further develop the RUS interface for more sophisticated use cases. Joint teleconferences between the UR and RUS working groups were also suggested to facilitate a close collaboration between the two groups.

A roadmap [8] was drawn up by Morris Riedel showing the new plans for the development of the RUS specifications that were generated during the OGF19 meeting.

All in all, OMII-Europe contributed enormously to the revitalization of the RUS and UR groups. Since OMII-Europe efforts, new communities (NGS, Brunel, DEISA) became interested in the standardization efforts of RUS and UR groups and thus the work within OGF is now well supported and covers a wider range than just the OMII-Europe project and its middlewares.

## **7 Outlook for Project Year 2**

The major milestone for the second year of the project is the release of the augmented components into the repository in December 2007 (M20). Table 3 shows a detailed break down of the tasks planned for each middleware.

During year two the main task for the DGAS is the development of the server component to support publishing and exposing Usage Record information contained in the legacy database.

The main focus of the SGAS development will be on the client side components to publish and query usage information via the new RUS interface.

The work within UNICORE 6 and its RUS interface and Usage Record generators will basically improve the alpha implementation towards production quality in year 2. Furthermore, a collaboration with the German national D-Mon monitoring project is planned that will start in June

2007. This will improve the resource level monitoring via the RUS interface and ensures that end-user communities of the German national Grid infrastructure D-Grid will have influence of the RUS interface development within UNICORE via feedback. Finally, some plans aim at a broader range of usage record generators that are currently only implemented for IBM LoadLeveler, Torque, LSF and PBSPro.

Middleware	Task	Finished by
DGAS	User clients and sensor ready for release	M16
	Standalone RUS server implementation	M16
	Synchronizer between RUS and legacy DB	M18
	Integration of sever and synchronizer into DGAS-HLR	M19
	Release of the integrated server component into repository	M20
SGAS	JARM (UR generation component) augmented with new URF	M13
	Integration of SGAS into ETICS build and test system	M15
	LUTS (server component) ready for release	M15
	JARM augmented with RUS interface ready for release	M16
	Augmented user client tools ready for release	M18
	Packaging for release and final testing	M20
UNICORE	Unit tests for the alpha implementation of the RUS interface	M16
	Integration of UNICORE 6 RUS implementation into ETICS	M17
	Interoperability tests within JRA3 – Task 2 with SGAS/DGAS	M18
	Matured version with Feedback from D-Mon Project & JRA3-T2	M19
	Packaging for release and final testing	M20

*Table 3: Planning for the 2<sup>nd</sup> year of the JRA1 accounting task.*

### **Planned Interactions**

It is planned to analyse the existing security solutions and make any necessary changes to the accounting components to meet the requirements of the common security profile from JRA3 task 1 during the second year. In conjunction with this adoption, feedback will be provided to JRA3.

To assist the JRA3 - Task 2 infrastructure integration, the accounting activity will provide semi-permanent service endpoints of the RUS server components to allow specification compliance and interoperability testing. The services will only be used for testing purposes and will not be maintained after the end of the project.

In conjunction with the release of the components the activity will coordinate itself with the SA2 quality assurance activity to ensure that the required quality criteria for release into the OMII-Europe repository can be met.

The activity is also planning to cooperate with the NA3 training activity to establish a module on grid accounting within the OMII-Europe courses.

Furthermore the possibility of installing the accounting components on the evaluation infrastructure is being considered in cooperation with the SA3 support activity. This especially applies to the modified components after release into the repository.

## References

- [1] Sven van den Berghe, Gilbert Netzer, Rosario Piro, Morris Riedel and Davy Virdee, JRA1 – Accounting Milestone: Preliminary design documents, 2006,  
<http://www.omii-europe.org/OMII-Europe/News/OMII-Europe-MJRA1.4.pdf>
- [2] Gilbert Netzer, Michele Pace and Morris Riedel, JRA1 - Accounting Milestone: Feature Complete Alpha Implementations, 2007
- [3] Annex I - Description of Work, 2006
- [4] J. Ainsworth, S. Newhouse and J. MacLaren, Resource Usage Service (RUS) based on WS-I Basic Profile 1.0, 2006,  
[https://forge.gridforum.org/sf/docman/do/downloadDocument/projects.rus-wg/docman.root.documents.version\\_1\\_0.draft\\_wsi\\_rus\\_17/doc7965](https://forge.gridforum.org/sf/docman/do/downloadDocument/projects.rus-wg/docman.root.documents.version_1_0.draft_wsi_rus_17/doc7965)
- [5] Rosario M. Piro, Evaluation and Design Plan of a RUS Interface for DGAS,  
[http://www.to.infn.it/grid/accounting/techrep/RUSandUR4DGAS-0\\_2.pdf](http://www.to.infn.it/grid/accounting/techrep/RUSandUR4DGAS-0_2.pdf)
- [6] Gilbert Netzer, Implementation Plan for the RUS Interface for the SGAS,  
<http://omii-europe.com/wiki/Wiki.jsp?page=SGASRusIf>
- [7] R. Mach, R. Lepro-Metz and S. Jackson, Usage Record - Format Recommendation, 2006,  
<http://www.ogf.org/documents/GFD.98.pdf>
- [8] Morris Riedel, OGSA - Resource Usage Service (RUS) Roadmap, 2007,  
<https://forge.gridforum.org/sf/docman/do/downloadDocument/projects.rus-wg/docman.root.roadmap/doc14261>