



EGI-InSPIRE

D7.2 ANNUAL REPORT ON OPERATIONAL TOOL MAINTENANCE AND DEVELOPMENT ACTIVITY

EU DELIVERABLE: D7.2

Document identifier:	EGI-docid-V9.doc
Date:	04/06/2012
Activity:	JRA1
Lead Partner:	INFN
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/1063

Abstract

The second year activity of the project work package 7 is reported in this document together with a spent effort assessment for all the tools. The roadmap for the next year of each tool is also mentioned but not elaborated in details since it was the focus of a recent project milestone.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Daniele Cesini	INFN	30/04/2011
Reviewed by	Moderator: Michel Drescher Reviewers: Peter Solagna Pavel Vaseanovici	SA2/EGI.eu SA1/EGI.eu RENAM	02/05/2012
Approved by	AMB & PMB		04/06/2012

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	23/04/12	First draft	D.Cesini
2	25/04/12	Internal review	D. Meredith/J.Casson
3	26/04/12	Internal Review	M.Babik

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:

<https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/about/glossary/>.



VII. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed – both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting ‘grids’ of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today’s production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.



The production infrastructure supports Virtual Research Communities (VRCs) – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

VIII. EXECUTIVE SUMMARY

This document reports the second year activity of the project work package 7 (JRA1). WP7 deals with the maintenance and development of operational tools. JRA1 is composed of five tasks:

1. TJRA1.1 is the management task
2. TJRA1.2 for the maintenance and underline developments for all the tools
3. TJRA1.3 devoted to the development of regionalised tools
4. TJRA1.4 for the extension of the accounting system to encompass new resource types (other than CPU)
5. TJR1.5 for the development needed to the extension of the Operations Portal and its harmonization with other portal frameworks

The second year is, for JRA1, the only one with all its tasks active, in fact TJRA1.3 should have ended at the end of PY1 but was extend for another year because of underspending and to complete its activities, while TJRA1.4 started in PY2 as planned.

TJRA1.2 activities proceeded regularly during the year and many new features were added to all the tools, these are described in details in section 2 and among them we can mention in this executive summary:

- **GOCDB:**
 - a new finer grained roles schema,
 - the implementation of service groups (also known as Virtual Sites)
 - the data scoping;
- **SAM:**
 - the development of a new profile manager (POEM)
 - the integration of the probes for all the middleware types currently used in production;
- **GGUS:**
 - a brand new report generator
- **Accounting system:**
 - a complete redesign for the accounting portal
 - a new data transport mechanism accompanied by a redesign of the accounting repository database
- **Metrics Portal:**
 - a complete redesign of the tool
- **Operations Portal:**
 - new modules added such as the **Security Dashboard** and the **VO-oriented Dashboard**.

For what concerns the integration of new (other than gLite) middleware types into the operational tools, this is now complete from the perspective of SAM and GOCDB for:

- UNICORE
- ARC
- GLOBUS
- Desktop Grids.



During the second project year, developments for regionalisation (TJRA1.3) proceeded slowly but the requirements were defined and analysed clearly and some progress in the development have been made, in particular for what concerns GOCDB with the data scoping that allows to distinguish non EGI sites.

The regional package for the Operations Portal was also improved during this year, however a full regionalisation for all the tools that were planned to be available for regional deployment is not available yet and will be carried on under the umbrella of TJRA1.2 in the coming years.

TJRA1.4 started in PY2 and, as a first step, a survey on what is available in various projects and standardization bodies for the accounting of resource types different from the CPU usage was carried on. Focus was given to the following type of accounting:

- storage accounting and the definition of its usage record
- accounting of virtual machines usage
- accounting of parallel and MPI jobs.

Some developments were done for the accounting of cloud resources usage within the project FedCloud task force.

TJRA1.5 activities are reported in section 2, in the paragraph dedicated to the Operations Portal.

A series of main issues arose during the year and are reported in section 5, we mention here the missing effort to face new big requirements for the SAM and Operations Portal teams, the second level support for regionalised tools that is done on a voluntary basis by too few people, the need for an inter-projects governing body for GGUS and the lack of a coordinated integration testbed for all the tools.

Section 5 also tries to provide a spent effort assessment for all the tools giving an estimate of the splitting between the effort needed for base maintenance and for new features development. This is done in order to give guidelines for the funding of future activities that will deal with operational tools development. The maintenance effort tends to increase for almost all the tools and the rise is steeper for newer or completely redesigned tools such as the metrics portal (all the effort devoted to development in the first phase for these tools).



TABLE OF CONTENTS

1 INTRODUCTION	7
2 DEVELOPMENT AND MAINTENANCE	8
2.1 Operations Portal	8
2.2 EGI Helpdesk (GGUS)	9
2.3 Grid Configuration Database (GOCDB)	10
2.4 Accounting Repository	11
2.5 Accounting Portal	13
2.6 Service Availability monitoring (SAM)	13
2.7 Metrics Portal	15
2.8 Broker network configuration and infrastructure	15
2.9 Integration of new middleware types	16
3 ACCOUNTING FOR DIFFERENT RESOURCE TYPES	17
3.1 Clouds	17
3.2 Parallel Jobs	17
3.3 Storage	18
4 REGIONALISATION STATUS	19
4.1 SAM	19
4.2 GOCDB	20
4.3 GGUS	21
4.4 Operations Portal	21
4.5 Accounting Repository and Portal	22
5 EFFORT ASSESSMENT AND MAIN ISSUES	23
5.1 Effort assessment	23
5.1.1 Development versus Maintenance Effort	24
5.2 Main issues	27
6 ROADMAP	30
7 PARTICIPATION TO CONFERENCES AND PROJECT TASK FORCES	34
8 CONCLUSION	35
9 REFERENCES	37
10 APPENDIX A1	40

1 INTRODUCTION

This document reports the second project year activity of the project work package 7 (JRA1). WP7 deals with the maintenance and development of operational tools. JRA1 is composed of five tasks:

1. TJRA1.1 is the management task
2. TJRA1.2 for the maintenance and underline developments for all the tools
3. TJRA1.3 devoted to the development of regionalised tools
4. TJRA1.4 for the extension of the accounting system to encompass new resource types (other than CPU)
5. TJR1.5 for the extension of the Operations Portal and its harmonization with other portal frameworks

The second year is, for JRA1, the only one with all its tasks active, in fact TJRA1.3 should have ended at the end of PY1 but was extend for another year because of underspending and to complete its activities, while TJRA1.4 started in PY2 as planned.

In this deliverable there is also the attempt to analyse and asses the effort used for each tool, trying to distinguish the effort spent in maintaining the tools and the effort used to develop new features. This is done, and will be probably repeated in the coming years, in order to give guidelines for funding the future activities/projects that will maintain the operational tools. The PY3 roadmap for the tools is not described in details in this document; it's just briefly outlined because it was the focus of a recent project milestone (MS708). A series of issues have arisen during the year and are also reported.

The document is organised into seven main sections:

- section 2 refers to task TJRA1.2 and describes the maintenance activity and the underline development needed to address the main requirements for all the tools;
- section 3 discusses about the extended TJRA1.3 and the status of regionalised tools;
- section 4 provides details about TJRA1.4 and the accounting for new resource types;
- section 5 reports the spent effort assessment and the main issues encountered during the year;
- section 6 is a quick overview of the next year roadmap for all the tools;
- section 7 reports the participation of the JRA1 staff to the main project conferences, working groups and task forces;

The document finishes with conclusions discussed in section 8.



2 DEVELOPMENT AND MAINTENANCE

This section gives a summary of the development and maintenance activity performed for each of the tools in the JRA1 scope during the second year of the project. A very short description of the tools is also provided at the beginning of each section together with all the needed references to get more information on tools architecture, documentation, deployment and future developments.

The general entry point to get information about the JRA1 activity and its tools is hosted by the EGI wiki at [R1].

2.1 Operations Portal

The Operations Portal is providing information to various actors (NGI Operations Centres, VO managers, etc.) along with related facilities, such as the VO administration tool, the broadcast and downtime system and access to the different specific dashboards (e.g. Operations dashboard, security dashboard). The Operations Portal is available as a central service and a regional package to be installed at the NGIs which will automatically synchronize information with the central instance. A detailed description of the whole system and of its modules can be found in [R2].

The development of the Operation Portal during the entire PY2 was focused on the following components

- 1) **The security dashboard** allows to collect results from Pakiti [R3] and Nagios [R4] in order to identify security issues, to display on a NGI basis the relevant sites with identified security problems, to create tickets against sites, to transmit results so as to prevent leakage, to compute security metrics based on the numbers gathered
- 2) **The VO management module and the VO ID card:** the VO management module allows to manage the VO ID card structure and contents and to check the VO ID cards registration and update using a persistent VO IDcard repository. The VO Management module is now used by COO (VO management is now operated by SA1).
- 3) **The VO oriented dashboard** allows showing aggregated status information of sites from test performed using VO Nagios systems. This implementation of this module was one of the main new features approved by the OTAG during PY2 for this tool.

A first prototype of **the security dashboard** has been delivered in July2011 and after evaluation from the EGI CSIRT group [R5] it was refined during PQ6, in particular for what concerns authentication (based on GOCDB and EGI SSO to allow restricted access to the data according to scope-based policies), visualization and metrics generation (customisable, dynamically generated and available as chart or tables in various formats). Various iterations with the CSIRT for the evaluation allowed improving the tool on each Operation portal release.

Concerning the **VO module** a number of small requirements coming from operations and user support teams were addressed in particular to provide more reliable statistics on users numbers and VOMS server information (to validate a VOMS host a "Test URL" button has been added to directly query the "listMembers").



The **VO oriented dashboard** development was approved at the OTAG-10 held during the Technical forum in Lyon and a first prototype was made available for evaluation in February 2012 and improved to address community comments in March and April.

The maintenance and requirements addressing for other components (i.e. Dashboard, COD Dashboard, Broadcast tool) continued in the limit of the available effort.

The **Regional Package** of the portal was maintained as well and the generation of the package (and the related documentation) is now done automatically. The synchronization with the central instance has been improved to authorize special characters in the messages. The broadcast module has been added in the regional version.

Decommission of the former CIC Portal had to be performed once all its feature were successfully migrated to the new Operations Portal and the last remaining feature (the User tracking tool) was migrated during PQ5, but some users were still relying on the old xml feed for the VO information available in the old instance, consequently the decommissioning will finally take place in PQ8.

During PY2 a total of 10 releases were performed by the Operation portal team, detailed release notes are available at [R 6].

2.2 EGI Helpdesk (GGUS)

GGUS is the EGI helpdesk [R7] and the major point of contact for EGI users to report service requests and problems. It hosts all project-wide support units and acts as the central integration platform in the distributed support infrastructure consisting of regional and topical helpdesk systems. GGUS is linked into various operational processes through interfaces to other operations tools.

There have been several important topics during the second year:

- 1) The continuous **integration of NGIs** into the support infrastructure. Thirty-three NGIs support units were created and integrated into GGUS by the end of PY2.
- 2) The refinement of the **Technology Helpdesk**, this is a module of the main tool requested by the project to handle all middleware related issues, and acts as a bridge between EGI and its external middleware providers. It was the major development task during PY2. Currently two workflows are handled in the Technology Helpdesk: (i) the software support workflow for bugs discovered in production; and (ii), the software provisioning workflow. In the latter case, this starts with a release announcement by one of the technology providers, and ends with the acceptance or rejection of the proposed software products.
- 3) The development and adaption work to have an **active-active fail-over system** for the data, the logic and the presentation layers of GGUS (Figure 1). This work will be continued in PY3.

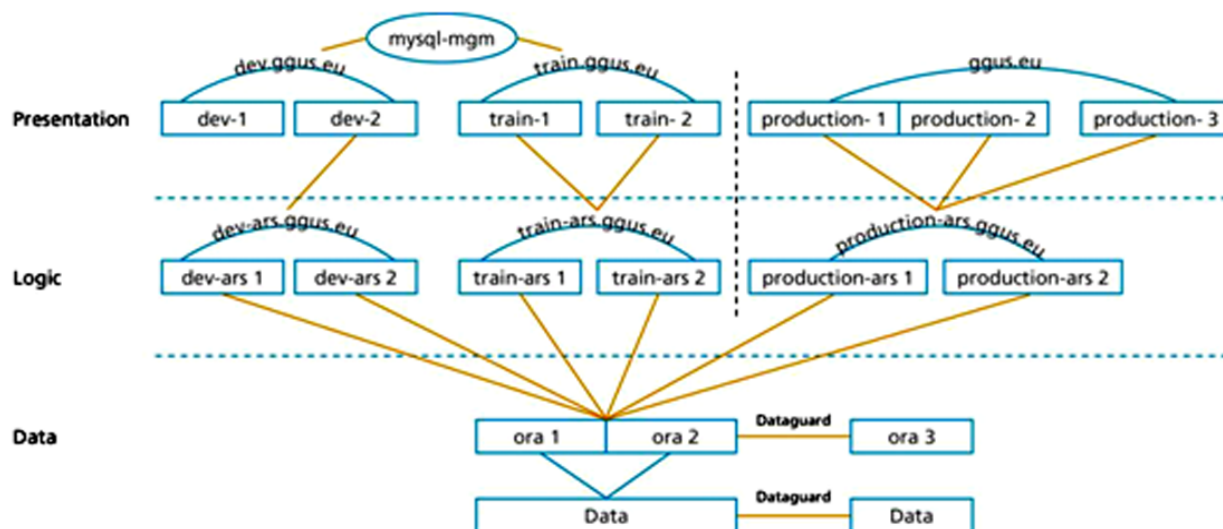


Figure 1: Schematic view of the high availability set-up of GGUS

During the first half of PY2 many requirements were created for the GGUS **report generator**, a module that allows to create statistics about tickets, support units and reaction times, so during the OTAG-10 face to face meeting in Lyon it was agreed to organise a workshop to analyse the requirements and propose a unique formalised set of requested features that could be implemented taking into account the available effort for GGUS within JRA1. The workshop was held on October, 26th and 27th at KIT in Karlsruhe, Germany and the agreed new features were:

- SLA reporting for EMI, IGE, SAGA
- Calculation of response times and solution times
- Calculation of statistical values like median, average, maximum, minimum of the data collected
- Reports by different parameters like priority, status, support unit, site etc.
- Ticket lifetime for user's view

A first prototype of the report generator was presented at the EGI Community Forum in Munich on March 2012 and the first usable version will be made available in June 2012 to have it finalised by the end of the year.

The regional instance for GGUS (xGGUS) will be discussed in section 4.3.

The GGUS product team published a total of nine releases during PY2.

2.3 Grid Configuration Database (GOCDB)

The GOCDB is the main source of topological information for the EGI Grid. It hosts general and semi-static information about the sites participating in the production infrastructure (covering data such as available services, service endpoints, service downtimes, site manager contact details, security contacts, virtual sites etc).



During PY2 effort focused on the three main activity streams: data model and backend, User Interface, Deployment. In the following some details are given about the developments for each of the mainstreams.

- **Data Model and Backend**

- a. Data scoping in the central GOCDB implemented in GOCDB-4.2, released 25-11-2011. The business logic, user interface and PI queries were updated to include new 'scope' tags used to label entities as belonging to the 'EGI' or 'Local' groups (similar in concept to a tag cloud). For more information see section 4.2 regarding a Regional-Publishing GOCDB and [R8]
- b. Service Groups (also known as Virtual Sites), which form logical groupings of services that are potentially distributed across a number of different sites and regions. See [R9]
- c. Refactorization of the backend PROM database to enable transaction demarcation in higher level business routines. Released in v4.1 on 01-11-2011 (v4.0 could leave the db in an inconsistent state if an error occurred since full rollback over lower level functions was not atomic)
- d. Update of wiki and documentation that now provides detailed information about ongoing developments, release notes and user guides
- e. New services and middleware types integration such as QoS-CoSGrid [R10] middleware and iRODS [R11] services
- f. Certification Status changes recorded in new audit/history table with accompanying PI queries
- g. Improved coding abstractions and integration of established design patterns and unit tests

- **User Interface**

- a. An improved user interface and a new Model View Controller logic
- b. Integration of a new XML output module based on the Query2XML package to generate complex nested XML documents from GOCDB data. This was required to render the Service Groups in the new 'get_service_group' PI method
- c. A new finer grained roles and permissions model. Released April 10th 2012. More information available at [R 12]
- d. Increased frontend responsiveness using db connection pooling

- **Deployment**

- a. A detailed versioning schema that was added to the tool releases starting from GOCDB v4.1 on 01-11-2011
- b. Established a GOCDB failover instance (http://bit.ly/frau_gocdb). The failover imports a secure download of the GOCDB data every 2 hours. A manually triggered DNS switch for the 'goc.egi.eu' domain between the production server and the failover server is in place. During normal operation, the failover is read only in order to prevent data-synchronization problems.

2.4 Accounting Repository

The EGI accounting repository [R13] stores information relating the usage of resources within the EGI production infrastructure. It receives data on individual jobs and summaries of collections of jobs

records from information providers, sites and other infrastructures, and exports accounting information to consumers of usage records, for example the accounting portal for visualization (see section **Error! Reference source not found.**). This is shown schematically in Figure 2.

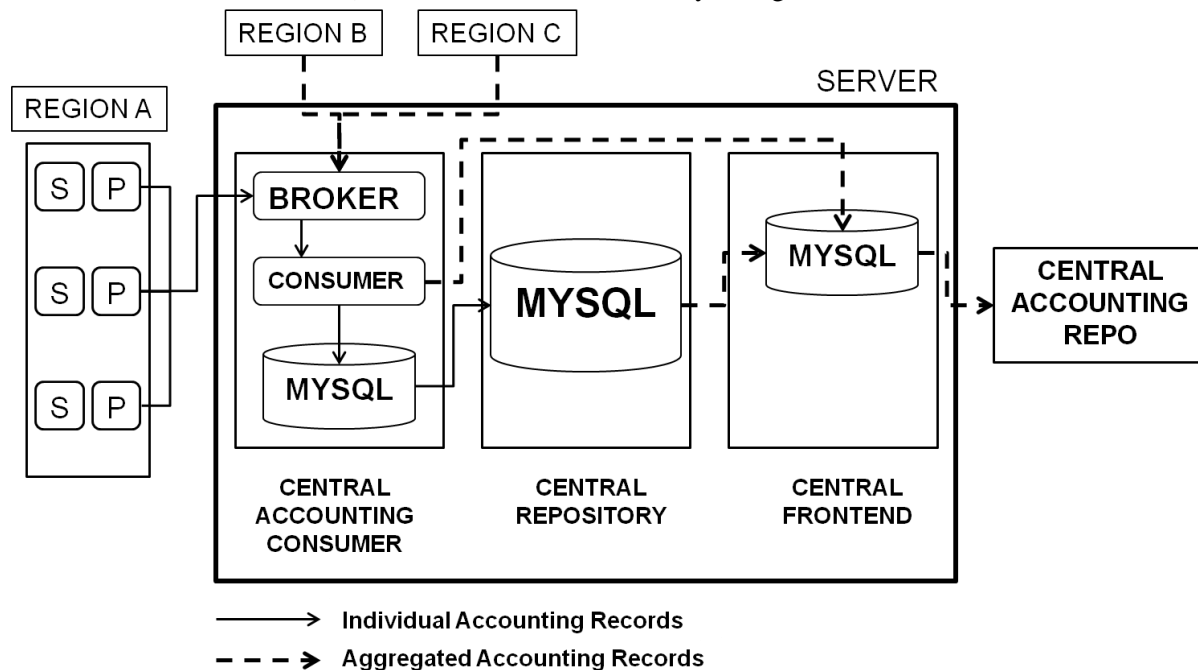


Figure 2: Schematic view of the Accounting Repository

As well as ingestion of accounting data from sites, the repository takes authorisation information from GOCDB and publishes the results of metrics to SAM.

Focus of the accounting development team during PY2 was, besides the maintenance of the production quality tool, on the migration of the system to a new transport method based on the Secure Stomp Messenger (SSM) which uses the production EGI Messaging Infrastructure. SSM transfers files in a secure and reliable manner between sites. SSM acts as a producer and consumer, taking files from a source directory at the publishing site, signing them, encrypting them, and sending them as messages to the receiving site where they are unencrypted and placed in a target directory. The necessary handshaking is also done using the messaging infrastructure to confirm the delivery.

A test service for the new STOMP/Python publisher/consumer was implemented in PQ5. Discussions with a variety of stakeholders including those that publish by direct database insert were started and continued during the entire PQ6. A training workshop was organised during the TF where the current accounting architecture and the planned roadmap were presented. Details and minutes of this workshop can be found at: http://bit.ly/tf_11.

In PQ7 and PQ8 the APEL team completed the APEL Accounting Repository major redesign and the development for the new transport layer (SSM) was completed; the new architecture is depicted in Figure 3.

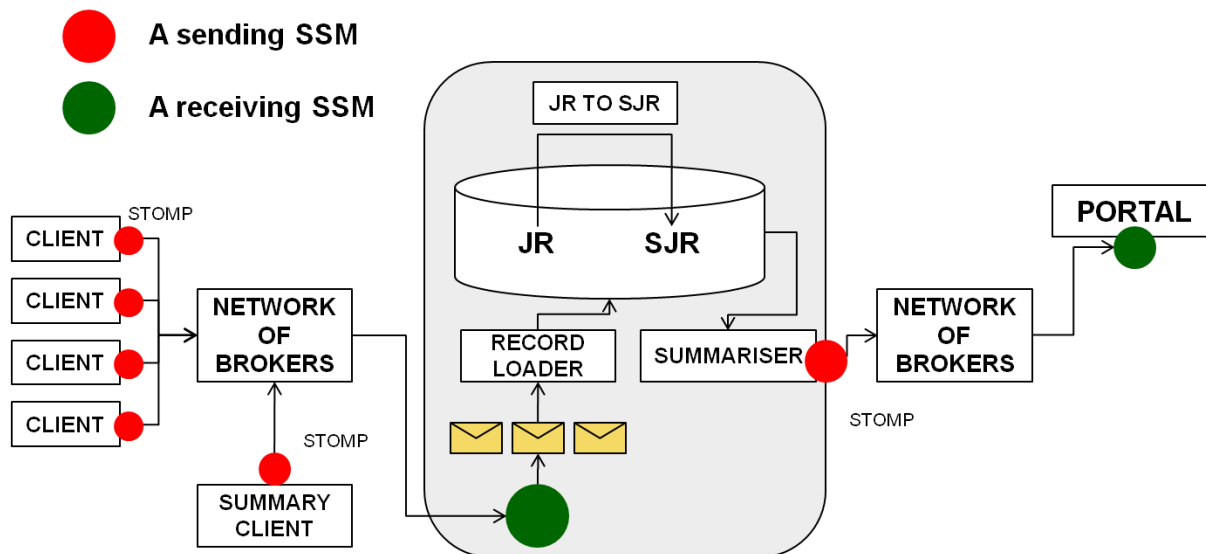


Figure 3: Outline of the new design of the Accounting Repository

2.5 Accounting Portal

The data recorded in the Accounting Repository (paragraph 2.4) are processed, summarized and displayed in the Accounting Portal, which acts as a common interface to the different accounting record providers and presents a homogeneous view of the data gathered and a user-friendly access to understanding resource utilisation. During PY2 the focus of the Accounting Portal team was on the redesign of the tool that now adopts a new codebase, a new graph engine and a new interface towards the Accounting Repository that now is based on the production Message Infrastructure. A query optimization was performed in order to increase the performance.

Many bug fixes and some new features were released to production. Among the new released features there are:

- Implementation of an XML interface to obtain data from the Custom View as requested by the user community
- Change of the Operations Portal interface from Oracle queries to XML (needed because the old Oracle interface was abandoned - the VO manager and user information would have stopped working)
- FQAN data on “VO admin” and “Site Admin” views as requested by the user community

The release process was also improved on a six months release base as described in [R14].

2.6 Service Availability monitoring (SAM)

The Service Availability Monitor (SAM) [15] is the system that is used to monitor EGI resources within the production infrastructure. It consists of the following components:

- Probes: a test execution framework (based on the open source monitoring framework Nagios) and the Nagios Configuration Generator (NCG)

- Aggregated Topology Provider (ATP), Profile Management System (POEM) and Metrics Results Store (MRS)
- a visualization portal (MyEGI)

EGI is not responsible for the probes development but only for their integration, the probe development hand-over to EMI (and other middleware providers such as EDGI and QosCosGrid) has been agreed and work is ongoing on the integration of the EMI probes with the SAM framework.

Figure 4 shows a schematic view of the SAM framework architecture.

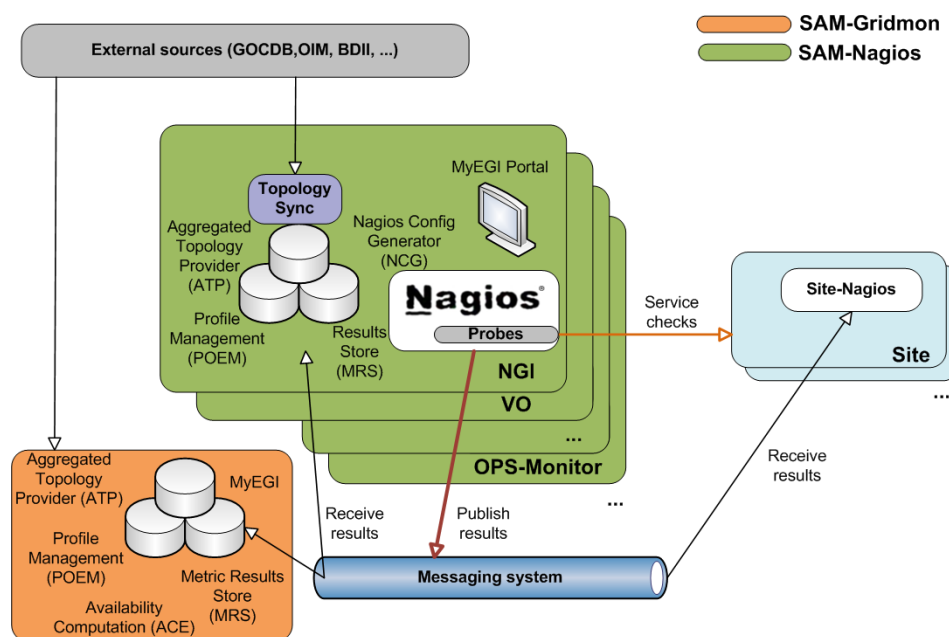


Figure 4: SAM components and their dependencies

A total of eight SAM updates were released to production during PY2, 3 in PQ5, 3 in PQ6, 1 in PQ7 and 1 in PQ8. All updates went through the full EGI Software Provisioning process¹ as was decided at the beginning of the project. The main activity streamlines on which the development focused are: 1) the inclusion of new middleware types probes 2) the improvement of the user interface for myEGI 3) the creation of a profile management system (POEM) to ease the addition of custom probes to the framework.

Concerning the inclusion of new middleware types probes, now the SAM system contains probes for gLite, ARC, UNICORE and GLOBUS. Desktopgrid probes are under integration. Other important probes that were included/improved in the system are those for GLEXEC, for the Certification Authorities and for other EGI operational tools such as GOCDB, GGUS, Accounting Portal and Metrics Portal. Moreover some of the internal SAM components can now be monitored (i.e. the MRS).

A first integrated version of the Profile Management system (POEM) was released in April 2012. The new profile management system introduced a major rework of the existing SAM architecture as it

¹ https://wiki.egi.eu/wiki/EGI_Software_Provisioning



completely changed the bootstrapping of all the major components and introduced significant changes to the overall data flow.

The MyEGI interface was improved during the year and many features were added, in particular all the features of the old gridmap [R 16] tool are now available in myEGI including the possibility to show the summary of Logical CPUs, Physical CPUs and HSPEC06 numbers for a given site.

Another important feature that was added in PQ6 is the possibility to configure a failover instance in a hot stand-by mode, more information available at [R17].

To ease the packaging and distribution of the system a clean-up of dependencies and metapackages reorganisation were also performed.

Decommissioning of the entire old SAM infrastructure including Gridview, Gridmap and the old SAM database and Web interface was finalized in April 2012 following successful deployment and transition to the new SAM/Nagios infrastructure.

2.7 Metrics Portal

The Metrics Portal displays a set of metrics that are used to monitor the performance of the infrastructure and of the project and to track their changes over time. The portal collects automatically through a series of connectors the relevant data that are published by the middleware or by other operations tools (i.e. GGUS tickets statistics) while some metrics are inserted manually by the activity leaders or by the NGI managers.

The metrics portal was completely redesigned with respect to the one available in the EGEE-III project, but it suffered hiring problems at CESGA during the first year of EGI-InSPIRE and the definition of its roadmap and development plan started and the end of PY1 after a series of meetings with the project activity leaders, held to get the list of requirements that had driven the development [R18].

At the beginning of PY2 effort was focused on the development of connectors to GGUS, BDII, and availability metrics data sources. The development portal was cloned to the production one. The possibility to manually override automatic metrics was added. In PQ6 the authorization model was redesigned according to the provided requirements and various enhancements were added (i.e. aggregated metrics on NGIs)

The following features were added during PQ7:

- Per country metrics (only on requested NGIs).
- Heavy query optimization.
- XLS output support.
- Many metric fixes and modifications.

A second major update (v2) for the portal was released in PQ8 enhancing all the portal components.

2.8 Broker network configuration and infrastructure

The JRA1 activity is responsible for providing support to the configuration of the message broker network of the production infrastructure based on the ActiveMQ system [R 19]. The focus for this activity during PY2 was on improving the authentication and authorization system to access the EGI production broker. Various updates of the ActiveMQ version were also performed. A system to enforce idle connection eviction was put in place in order to greatly improve the performance of the



system. A plan for the implementation of a credential-sharing procedure between brokers that are members of the same network was also created. This is needed in order to have the same credentials available to all brokers and allow clients to choose on their own which endpoint to use. This task has to be finished before requesting clients to migrate to authenticated connections.

A lot of effort was spent to support migration from “topic” to “queues” (that enables the message history retention) of other operational tools, in particular the Operations Portal, allowing for a more reliable message exchange between the tools.

2.9 Integration of new middleware types

Integration of new middleware and resource types affects almost all the operational tools, but in particular the GOCDB that has to record the inclusion in the production infrastructure of the new resources, the SAM framework that has to monitor them and the accounting system (both repository and portal) that has to provide accounting information.

At the time of writing GOCDB contains 66 services types from the following middleware stacks: gLite, UNICORE, Globus, iRODS, ARC, QosCosGrid. In addition gateways to Desktop Grids are also recorded. A full list of services types can be obtained from the GOCDB programmatic interface [R20]. Inclusion of new service types in the GOCDB requires the central operators’ intervention and in the last project quarter, a procedure was defined for service-inclusion that requires the OTAG approval. This is to avoid uncontrolled insertion and duplication of service types. For custom service types that are not recognized as generic middleware/service components, a new naming scheme was agreed which involves pre-fixing the service type name with the ‘CUSTOM.’ prefix. This facilitates a lighter weight and more streamlined procedure for the addition of custom services.

As reported in section 2.6 the SAM monitoring framework is now able to monitor services from the following middleware stacks: gLite, UNICORE, Globus, ARC and Desktop Grids.

Operational tools can be considered as new services types and probes have been implemented and integrated in the last SAM update in order to check their availability.

Inclusion of new middleware and resource types in the accounting system is discussed in section 3.

GGUS is also indirectly touched by the inclusion of new middlewares in the production infrastructure, in particular for what concerns the support units to be added to the technology helpdesk that now can handle tickets for: EGI DMSU [R21], EMI [R22], IGE [R23] and SAGA[R24].

Concerning the tools to operate the federated cloud infrastructure, many JRA1 representatives attended and actively participated in the EGI FedCloud task force activities [R25]. JRA1 led two working groups (also called Scenarios), one for monitoring (Scenario 5) and one (Scenario 4) for implementing an accounting system of the cloud infrastructure. Within Scenario5, JRA1 provided a NAGIOS box and the integration of the cloud resources probes developed by some of the partners participating to the task force. The FedCloud monitoring system is now in place [R26]. Scenario4 activities are detailed in Section3.1 of this document.

3 ACCOUNTING FOR DIFFERENT RESOURCE TYPES

One of the JRA1 objectives is the evolution of the accounting system to encompass the new different resource types that are joining the production infrastructure. Moreover JRA1 will evolve the accounting system to support an economic model needed for the self sustainability of EGI.

The task responsible for the needed development is TJRA1.4 – Accounting for different resource types – which is a 3 year task started at the beginning of PY2.

The activity for TJRA1.4 during the firsts two quarters of PY2 focused on the investigation of what is currently available for the accounting of new resource types, identifying overlaps and possible collaborations in the work performed by various stakeholders, trying to bring many of them together. This activity culminated in an accounting development workshop held in Lyon during the TF2011.

The types of accounting considered, and the stakeholders were:

- CPU (inc OGF UR) (EGI, EMI, OGF)
- MPI (EGI, EMI)
- Storage (inc StAR) (EMI, OGF, EGI)
- Virtualisation (EGI, other projects)
- Applications (EGI)
- Data Use (EU-DAT, PaNData)

The agenda and minutes of the workshop [R 27] [R28] report details about each of the previous points. The following paragraphs describe the activity performed on clouds, parallel jobs and storage accounting.

The JRA1 accounting product team is also participating to the following task forces:

- integration of new middleware types TF (i.e. UNICORE[R29])
- TCB inter projects Accounting TF [R30]

to get requirements about the integration of new middleware stacks (ARC, UNICORE, gLite, Globus).

3.1 Clouds

The JRA1 accounting product team actively attended the project FedCloud [25] taskforce and led the so-called “Scenario4- accounting” activities with the mandate to include into the accounting system records coming from the federated cloud infrastructure. A “straw man” Usage Record based on the CPU one has been defined and the partners running the various infrastructures are reporting back on how well they can cut them. Prototypal testing with cloud Usage Record following from the resource centres to the central accounting repository started before the EGI Community Forum in Munich and was shown during the FedCloud demo at the Forum. It will be improved in the following months and will be more reliable once the new production infrastructure based on SSM is in place.

3.2 Parallel Jobs

The JRA1 accounting product team also participated in the definition of the accounting for parallel jobs usage record within standardization bodies (i.e. OGF UR-WG). It was confirmed that it was sufficient for parallel jobs to publish at least the number of compute nodes used by a job [R38], the need for publishing also the number of physical CPUs is under debate. It is intended that CPU elements in a UR record contain the number of cores for compatibility with the accepted practise for



serial jobs running on a single core. Unfortunately APEL did not implement these fields so that is a target for 2012. The repository will be ready to receive them when the database is migrated but will not actually receive the data from most sites until the new APEL client is rolled out.

3.3 Storage

The JRA1 product team also participated in the definition of the accounting storage record within standardization bodies (i.e. OGF UR-WG). Meanwhile it is moving forward as an interim standard in EMI where the storage product teams are preparing to populate records. They will publish using SSM and the APEL repository will receive the records and load them into a database. The release of the EMI products is not scheduled until EMI-3 in May 2013 but this may be negotiable if EGI has a strong requirement



4 REGIONALISATION STATUS

The JRA1 task TJRA1.3 deals with tools regionalisation, it is intended to carry on the development needed to have regional, deployable versions of many of the JRA1 tools including GOCDB, GGUS, SAM, Operation Portal, Accounting Repository and Accounting Portal.

This task was originally intended to be active only during the first year of the project, but at the end of PY1 it resulted to be highly unspent from many partners and after the first EGI EC review it was decided to extend the task for another year. The task was unspent because the regionalisation development didn't progress for many of the tools for three main reasons: 1) the complex and tightly coupled interconnection among the tools that slow down the development of regional instances - since every delay in a tool implies delays on all the others and 2) the absence of well defined requirements for some of the regional tools 3) higher priority developments that have emerged over the period that affect the central instances. In example, for GOCDB this includes the new role/permissions model, virtual sites, data scoping, essential refactoring etc; these developments needed to be completed before regionalisation work could realistically commence. To face at least the second point, at the end of PY1 the OMB approved the creation of a dedicated taskforce with the purpose of analyzing the regionalized operational tools use cases and the dependencies between different regionalized tools. The output of the task force was a series of requirements to JRA1 for development that are summarised in the following for each of the tools together with the work that was performed by the developers to address the requirements.

4.1 SAM

SAM can be deployed as regional instances that synchronise with the central instance. SAM was already fully regionalised at the end of PY1 and is released following EGI software release process (including Staged Rollout). A list of participating NGIs can be found at [R31].

ATPs at NGI instances are independent from the central one, meaning that they perform synchronization with information sources directly. NGI instance could in theory point ATP to different information sources (e.g. alternative VO feeds or regional GOCDB).

The requirements coming out from the regionalised task force were driven by two use cases: 1) monitoring local sites and services not recorded into the GOCDB and 2) adding custom probes to the regional instance and using different VO for different services.

From a JRA1 perspective both requirements look valid but very complex to implement and given the amount of effort available it looks unfeasible in the short term to fully implement them.

Monitoring sites and services not recorded in GOCDB is a complex task as SAM would need to find alternative ways to receive site and service downtimes, contacts, service types as well as ensure this information is consistent and doesn't contradict what is already recorded in GOCDB. In other words it needs to supplement, to some extent, the functionality of GOCDB. Depending on the actual use cases, it might be also necessary to support topological and profile history in order to be able to provide clear explanations of the computed availability and reliability numbers.

Concerning the integration of custom probes into the regional instances, this can now be done and the release of the POEM component eases the new probes management; ad hoc documentation will be created to describe the needed steps, but this will only be possible for services and sites already recoded in GOCDB due to the previously stated reasons.

4.2 GOCDB

The GOCDB is mainly used as a central service. A standalone GOCDB is freely available for download, but it is not possible to automatically synchronize the data hosted in a local installation with the central instance. The regionalisation task force raised two main requirements:

- 1) A hierarchical architecture for the GOCDB so that NGIs can exploit at least two layers of the tool: regional GOCDBs populating the bottom layer, and the EGI central GOCDB installation populating the top layer. In this model, information would flow from the bottom layer to the top layer, with the possibility to filter data when sending to upper level and
- 2) Customization and extension of the GOCDB schema allowing NGIs to extend their GOCDB for internal/specific purposes that do not involve other NGIs. Importantly, those customisations should not be implemented in the central GOCDB instance.

Requirement 1) would allow the introduction of local sites in the regional infrastructure (regional nagios and the regional dashboard) without propagating that local site information to EGI beyond the NGI scope.

Both these requirements are difficult to address with the current effort available for the product team. Their implications were carefully analysed and a 3phase plan was created [R32]:

- 1) **Data Grouping/Scoping:** A GOCDB instance (whether central or regional) first needs to differentiate between EGI and non-EGI data. To do this, data scoping was introduced into the central instance allowing GOCDB data to be tagged as either 'EGI' or 'Local' in scope. Functionality to insert and remove data to and from a scope was also provided. With scoping in place, NGIs are able to store and retrieve both local and EGI scoped sites and services in the central GOCDB. This allows users to centrally include or exclude their sites and services from the EGI infrastructure. Scoping in the central instance therefore provides a single point to input and query data for both EGI and local/regional data. In doing this, a regional GOCDB installation is not strictly necessary for those NGIs without effort to install and maintain a regional GOCDB.
- 2) **Regional-Standalone GOCDB:** a freely downloadable standalone instance of the GOCDB will be maintained.
- 3) **Regional-Synchronised GOCDB:** this would allow NGIs to deploy their own GOCDB, customize their 'Local' scoped data, and publish only their 'EGI' scoped data to the central instance. Publishing of this data would most probably be via a schema constrained WS/REST interface. The central portal would remain as the definitive source for all EGI information.

The first step of the previous list was released to production during PQ7 with GOCDB-v4.2 and it was extensively documented in the EGI wiki [R33] and [R34].

The second step, the stand alone GOCDB, is already available and technical documentation is provided at [R35]. The biggest limitation of the currently available download is that it requires an Oracle database. Although an Oracle XE database is freely available from Oracle, migration to other databases should be prioritized by the OTAG as needed.

The last step, the regionalized-synchronized instance, is a long-term development. Given the available effort for this activity, it is currently not possible to provide a definitive timeline. In addition, the GOCDB product team has a number of concerns about this deployment model that cannot be trivially addressed. A much simpler solution would be to continue to use the central GOCDB to store EGI scoped data, while a regional GOCDB could present views and edit forms of the central GOCDB data (as inner frames) alongside its locally scoped NGI data. Moreover, another barrier for a regional-



publishing GOCDB is that it would be far less customizable compared to a regional-standalone GOCDB for the following reasons:

- Central updates will inevitably introduce new features (recent examples being new user roles and virtual sites). This will change fundamental GOCDB behavior and will almost certainly break compatibility with any NGI specific customizations. How to treat NGI customizations compatibility need to be carefully analyzed and regulated by OLAs before starting the development of a regional publishing instance
- For each GOCDB update, the NGIs will have to provide effort to test and fix their customizations.
- The central GOCDB team should not be involved in supporting bespoke NGI customizations.

Given these concerns about the regional-synchronized GOCDB, it is clear that JRA1 needs detailed discussions with all the involved parties (NGIs in particular) about this requirement. This will take place in one of the first PY3 OTAG meetings.

4.3 GGUS

GGUS was never intended to be a regionalized, deployable instance. The regional implementation since PY1 is provided by xGUS, which is currently a *custom service* hosted *centrally*. In xGUS the tickets can have a local or global scope, all answers to a global ticket are redirected to the central GGUS. Many NGIs are using their own support systems such as RT and they are, since a long time, already interfaced and synchronised with GGUS. The only use case that was raised at the regionalisation task force is about the possibility to automatically set up regional sites into xGUS. Currently xGUS doesn't sync with regional or central GOCDB. Sites have to be inserted by hand in the helpdesk configuration. A regional helpdesk should be able to be configured with a regionalized GOCDB in order to automatically insert sites, both EGI or local ones (i.e. if the central GOCDB with multiple views is implemented, global and regional views, xGUS should be able to import sites from both the views). The possibility to edit the sites list by hand should remain as a backup solution. This requirement was however agreed to be a low priority one.

This requirement is partly already in place since site support units can be imported via the xGUS web interface and the portal administrator can choose which registered sites should be imported. Prioritization of the development needed to address the rest of the requirement will take place during PY3.

4.4 Operations Portal

A first implementation of the regional package of the Operations Portal was released in June 2010 and it was continuously updated and maintained up to now. There are currently four instances of the regionalized operations portal used by: Portugal/Spain, Greece, Belarus and Czech Republic. The NGL_UK will set-up a regional instance in the first semester of 2012. The package and the documentation are available through a SVN server [R36].

The local instances push the information (alarms, tickets, metrics) to the central one. In case of disruption of services in the Central Instance the regional ones continue to operate properly. The main use cases envisaged by the regionalization task force are:

- 1) The possibility to switch on and off new alarms: as required by an NGI within the ROD scope only. Consequently, rod on duty can be alerted for all those alarms considered critical for his region.
- 2) The integration of local non-EGI sites and services into the dashboard: normal dashboard workflows should be extended to allow RODs to operate on both sites not listed in the central



GOCDDB and single services operated by EGI sites that are not in the GOCDDB either not in the set of services recognized by the EGI infrastructure.

3) The use of a different helpdesk for local sites

The proposal is to use the regional GOCDDB to solve use case 2), so the solution depends on that development as described in section 4.2. However currently the regional Operations Portal cannot use more than one GOCDDB, so all the services/sites both EGI and local, should be in the same view. The support of more than one GOCDDB view is a feature foreseen for the regional instance of the Operations Portal. For use case 3) the development team is already working on the solution interfacing xGUS.

4.5 Accounting Repository and Portal

The regional accounting repository is currently under development and its first release is expected to happen in December 2012, however it requires that the production infrastructure will be migrated to the new SSM based system. The regionalization of accounting repositories includes two scenarios:

- A local release of the central repository that a region can use. The model is hierarchical, clients will publish on the regional repositories that forwards the summary records to the central one.
- Regions run their own accounting services that need to be integrated with the new accounting server.

For NGIs that do not deploy a regional repository it will be possible to continue direct publishing from sites to the central repository.

The Accounting Portal does not need specific development for a local installation. The code for the regionalized instance is almost the same as the global instance. A regionalized accounting portal is basically a global portal that queries the data from a regionalized accounting repository.

The only requirement that came out from the regionalization task force concerns the possibility for a regional portal to publish centrally only data for international VOs keeping private those about national VOs. The requirement has been recorded in RT but up to now there was no available effort to address it.

5 EFFORT ASSESSMENT AND MAIN ISSUES

5.1 Effort assessment

Task composition of the activity was discussed in details in D7.1 [R37] and is summarized here:

- **TJRA1.2: Maintenance and development of the deployed operational tools:** The underlying bug fixing and development work for the operational tools
- **TJRA1.3: National Deployment Models:** Development needed to support the regionalisation of the tools at the NGI level
- **TJRA1.4: Accounting for different resource types:** The work needed to evolve the EGI accounting system in order to encompass the different resource types that will be included into the production infrastructure and to support an economic model needed for the self sustainability of the EGI Grid
- **TJRA1.5: Integrated Operations Portal:** Dedicated the development needed for a restructuring of the Operations Portal, its evolution and harmonization with other portal frameworks

The JRA1 tasks follow a particular time sequencing: TJRA1.3 was intended to last only for PY1, TJRA1.4 for 3 years starting from PY2 while TJRA1.5 for the first 3 years of the project. TJRA1.2 (and also the management task TJRA1.1) are the only tasks that last for the entire project duration. After the first EC review it was decided to extend TJRA1.3 for another year to face a strong underspending of the task and to complete the needed developments that didn't happen in PY1 as discussed in Section4. Table1 summarises the time sequencing of the JRA1 tasks, black box means task not active.

	PY1	PY2	PY3	PY4
TJRA1.1	Active	Active	Active	Active
TJRA1.2	Active	Active	Active	Active
TJRA1.3	Active	Active	Inactive	Inactive
TJRA1.4	Inactive	Active	Active	Active
TJRA1.5	Active	Active	Active	Inactive

Table 1 - JRA1 tasks time sequencing. TJRA1.3 was extended for another year at the end of PY1

Committed and used effort data for the development tasks, TJRA1.2, TJRA1.3, TJRA1.4 and TJRA1.5 are reported in the tables available in Appendix1. They show the data for the first year of the project, for the first three quarters of the second year and for the sum of the two periods when possible (some tasks could have been inactive in one of the two timeslots).

Task by task the following considerations apply:

TJRA1.2

The total TJRA1.2 effort consumption is in line with the committed effort, even if for some of the partners deviations arise. In particular GRNET underspent during PY1 and is still underspending on this task, while FCTSG/CSIC is overspending.

TJRA1.3

The strong underspending of the first year now is almost fixed in the total number (93%) but the some underspending remain for FCTSG and STFC due to the fact that the development for regional GOCDB and accounting system is not yet completed.

TJRA1.4

TJRA1.4 shows a quite strong underspending in the total numbers for the first three quarters of the second year (that is also the entire duration of the task since it was inactive during PY1). This underspending was expected since most of the work done in this initial phase consisted mainly in requirements collection and study of what is currently available in the various middleware/projects for the accounting for different resource types. It will be compensated during the next two years when a more active development phase will start.

TJRA1.5

No big deviations from the committed effort, just a slight underspending that is not due to any particular reason and that is mainly coming from the first year of the project. It will be probably compensated in the next year when the big requirements that were recently submitted to the Operations Portal will be addressed by a development phase. The new transversal requirements on to the Operations Portal will induce a larger workload which will be by no means visible in the (“committed and used effort”). Indeed, actual workload cannot be based for the Operations Portal on the claimed effort in PPT versus the PPT targets of the initial EGI-Inspire roadmap – see issue 9 in section 5.2.

5.1.1 Development versus Maintenance Effort

In order to assess the effort needed in the longer term and to give an idea to future funding sources on how to guarantee at least a minimal level of maintenance for the operational tools, it’s important to start evaluating which was the effort used tool by tool and how the used effort is split between basic maintenance (bug fixes, interaction with other tools) and development for new features or code restructuring etc. The evolution of both effort numbers during the entire duration of the project could gives at least a trend of the maintenance cost of the tools. Cost in this context is used in terms of effort and not of money, translating effort to money is outside the scope of this deliverable.

The project effort used tool by tool and the its splitting in maintenance versus development cannot be inferred directly from the tables presented and discussed in the previous paragraphs because:

- not all tasks are mapped one to one with the tools
- the same partner can develop more than one tool
- JRA1 does not have specific tasks dedicated only to maintenance.

Even if some tasks could, as first approximation, be considered development-only tasks (i.e. TRJA1.4), this is not true for other tasks such as TJRA1.2 which contains most of the effort for almost all the partners. The splitting should be self-assessed by the product teams for each of the tools. This assessment for the first two years of the project, tool by tool, together with a forecast for PY3 that will need to be validated during the next project year are reported in Table2, Table3 and Figure5. Table2 shows the absolute numbers, tool by tool, of the total project effort used (in PY1 and PY2), while Table3 reports the splitting between the development and maintenance.

The maintenance effort tends to increase for almost all the tools and the rise is steeper for newer or completely redesigned tools such as the metrics portal being constantly upgraded and improved at a high pace (all the effort devoted to development in the first phase for these tools).

The accounting portal case is different because of the very low effort that was assigned to it at the beginning of the project due to the hiring problems that affected the FCTSG/CSIS staff, so almost all the available effort was devoted to maintain the code inherited from the EGEE tool.

Note that with the end of TJRA1.3 almost all the effort available to the SAM team will go into maintenance, meaning that many new requirements for the tool could not be implemented during the next years, this will be reported as an issue in the next section.

The fraction of the maintenance effort for the Operations Portal was maintained low during PY1 and PY2 because of the high number of new and complex features requests that was addressed by the PT. This was done exploiting the TJRA1.5 effort (18PMs per year, for the first 3 years, dedicated only to the development of new features) and thanks to the additional manpower that the France NGI committed throughout the project duration (4 years). This additional manpower allowed to keep high the fraction of effort devoted to development. Also, absolute numbers of maintenance effort are expected to increase for the Operations Portal in the coming years because of the many newly added features, that in general are reflected into the creation of new sessions and modules of the Portal (see issue9).

For GOCDB, the balance between maintenance and development depends upon regionalisation. If the regional-publishing GOCDB will still be required, then development would need to be increased.

	Involved Partners	PY1	PY2(excluding PQ8)
GOCDB	STFC	5	5
Ops Portal	CNRS	17	17
GGUS	KIT-G	10	11
SAM (including messaging support)	CERN SRCE GRNET	13	10
Accounting Repository	STFC	1,6	13 (including tjra1.4)
Accounting Portal	FCTSG/CSIC	6	3
Metrics Portal	FCTSG/CSIC	0,6	1,76

Table 2 – Reported effort tool by tool for the first two years of the project, excluding PQ8

Tool	PY1		PY2		PY3 forecast	
	Dev	Maint	Dev	Maint	Dev	Maint
GOCDB	50%	50%	50%	50%	60%	40%
Ops Portal	80%	20%	80%	20%	70%	30%
GGUS	55%	45%	55%	45%	50%	50%
SAM	45%	55%	35%	65%	25%	75%
Accounting Repository	75%	25%	75%	25%	60%	40%
Accounting Portal	20%	80%	30%	70%	40%	60%
Metrics Portal	90%	10%	90%	10%	50%	50%

Table 3 - Maintenance versus Development Effort, tool by tool, for the first two years of the project and the forecast for the third.




Figure 5 - Evolution of the fraction of effort devoted to maintenance tool by tool for the first three year of the project. The point for PY3 is a forecast.

5.2 Main issues

The following main issues have been found during the second year of the project:

- 1) **Effort for SAM new developments:** The effort available to the SAM team for new development is very limited taking into account the complexity reached by the system, this is also made worse by the end of TJRA1.3 (3PM per year to the team). Some important requirements (i.e. support for virtual sites, monitoring non-EGI sites and services) are on hold, waiting for the needed effort to be available. All the effort will be mainly used for maintenance and 3rd level support.
Mitigation: Where possible new developments will be moved to other operational tools (i.e. availability and reliability calculation for virtual sites) and in order to limit the 3rd level support requests, the 2nd level staff should be regularly trained.
- 2) **Available effort for regionalisation activity is not enough:** As discussed in section 4 addressing some of the regionalisation requirements is very effort consuming. In particular the creation of a synchronised regional GOCDB and the monitoring of sites and services outside the EGI scope from SAM cannot be addressed in the short term and, even in the case that JRA1 will manage to carry on the needed developments, their maintenance in the long term will be very difficult without a significant increase in the available effort for the product team. Moreover effort for customized support of the regional package of the Operations Portal is increasing with the numbers of NGIs adopting the tool and the team has deliberately stopped regionalizing more features of the operations portal
Mitigation: An alternative solution has been proposed for the regional synchronised GOCDB that needs to be discussed within the advisory group, while for SAM apply the same considerations given in the previous point. Not all the newly developed features of the Operations Portal are ported to the regional version.
- 3) **Integration testbed missing:** because of effort reasons, JRA1 never managed to create a production-like environment for testing new updates of central tools. Each product team for their products does a lot of testing but what is really missing is an integration testbed to verify that the complex dependencies among the tools are not broken by new updates. Moreover real users and real-life use cases are strongly needed for testing. A recent example is given by the release of the new roles schema in the GOCDB that made some components of the Operations Portal unusable by certain GOCDB roles. The integration testbed should not be created only by JRA1 for a couple of reason: effort is missing and real users that can provide actual use cases are outside JRA1.
Mitigation: start discussing the issue with other actors (i.e. tools users) involved, mainly SA1 but also NA2/SA3 for the VO-oriented components of the tools.
- 4) **Staged Rollout for Operations Portal:** The Operation Portal is a regionalised tool since PQ1 even if it is not deployed in all NGIs, however it doesn't follow the standard EGI release process which involves quality criteria verification, staged rollout and inclusion into EGI official repositories. The experience gained with the SAM tool in the EGI software release mechanism is very positive and showed that the process avoided critical bugs to be released to production in some cases: the small delay introduced is compensated by the higher quality of the software, so it's interest of JRA1 to have all the regionalised tools released using the official software release channel. However, Operations Portal structure and complexity has

to be carefully explored to establish whether release process adopted for instance with SAM can be generalized to the Operations Portal, e.g. packages released should not disrupt regionalized view in the central version. Indeed it is not possible to operate in the same time a regional portal and the related regionalized view in the central portal.

Mitigation: Start discussing with SA2 the release process for Operations Portal taking into account that for the major development of the central portal, pilot versions are already tested by concerned users (e.g CSIRT group for the security dashboard, VO Managers for the VO dashboard)

- 5) **2nd level support for SAM still done in a voluntary basis:** while effort was originally allocated by the project for the 2nd level support of middleware services (i.e. DMSU) no effort was foreseen for operational tools. In particular for widely deployed tools (currently only SAM), this is a very time consuming activity that subtracts important resources to developments because it is done directly by the product teams. Even if not deployed at the same level of SAM also the 2nd level support for the Operation Portal is very time consuming. To mitigate this issue at the beginning of the second year of the project it was decided to instantiate a team of volunteers, chosen among NGI experts, to run a 2nd level support unit in GGUS for SAM. The team was trained remotely for the first months and then face-to-face during the EGI community forum in Munich. It performed an effective work unloading the product team from following many deployment problems and tickets. The problem is that JRA1 together with SA1 managed to find only three volunteers to create this support unit which is now becoming quite overloaded by this unfunded activity. Without a greater number of supporters, that would give the possibility to create a reasonable rota inside the team, the risk is that that the volunteers decide to stop participating to the activity.

Mitigation: New supporters should be found within the operations community and to favour this research a dedicated effort should be allocated by the project management. During the EC review JRA1 will propose to move some effort available at NGIs to the 2nd level support of regionalised operational tools, not only SAM

- 6) **GGUS governing body:** while most of the JRA1 tools are used mainly by actors belonging to EGI or closely linked to EGI (operations community, project management and VO managers) this is not true for GGUS which is used by all the Grid users and many of them could have weak relationship with EGI (i.e. not highly represented in EGI bodies, not having an EGI SSO account). The inclusion in GGUS of the technology dashboard, used to link users and developers (i.e. from EMI and IGE), also increased the number of GGUS users outside EGI. This means that the EGI governing bodies of operational tools (OMB and OTAG) often miss important representative of heavy user communities of GGUS and are not the most suitable place to discuss GGUS requirements prioritization. During the second year of the project JRA1 tried to held extended (to the relevant communities) OTAG meetings and when needed, for instance to discuss the new GGUS report generator, ad hoc task forces were created. This resulted in useful meetings but the frequency we managed to have was really too low compared to what is needed to face promptly the requirements coming for GGUS from the many actors involved. This requirements quite often do not imply big changes, but need to be addressed quickly (i.e. NGI and support units creation) and cannot wait for the organisation of huge extended meeting.

Mitigation: officially instantiate a permanent, cross projects, small group to discuss with high frequency GGUS requirements. The group already exists de facto, it includes JRA1 developers and WLCG representatives and should be officialised extending it to other relevant communities.

- 7) **GGUS bugs and requirements tracking tool:** for the very same reasons discussed in the previous point it was difficult to fully migrate the GGUS tracking tool for bugs and new



features requests from the CERN savannah used before EGI and the RT based system adopted by EGI to track requirements for middleware and operational tools. Some work was performed (within JRA1 and WLCG) in order to understand if all the features needed by the user communities and by the development team are available in the new system, but this work never completed. The authentication and authorization issues for people outside EGI were also analysed but an agreed proposal for a solution that could allow a smooth migration from savannah to RT for all the involved actors was not reached, so at the moment both systems are used to track GGUS requirements which is not an optimal situation in particular when the requirements have to be discussed for prioritization.

Mitigation: instantiate a inter projects task force to complete the work started during PY2 to reach a decision on the issue.

- 8) **Single access point to releases information:** as extensively discussed in D7.1 [R37] at the beginning of the project it was decided to organise the JRA1 activity as set of independent product teams with a central coordination. This means that each PT has its own bugs tracking tool, building system, code repository and documentation pages. The requirements tracking tool was unified according to the project guidelines (with the exception discussed in the previous point for GGUS) and the EGI wiki was used as a central access point to reach the distributed information for all the tools. Despite the efforts dedicated to maintain updated the wiki pages it is still sometime difficult to get information about the schedule and content of next releases. This information for each tool of course exists, but a single aggregation point that shows in just one page what is going on within JRA1 is still missing. We tried to use a dedicated RT queue, the “ops-tools-roadmap” queue, to provide this information but maintaining the information updated was really time consuming and resulted in a duplication of effort since also the internal PTs pages had to be updated with the same data. The usage of the “roadmap” queue is now abandoned.

Mitigation: a dedicated section with all the needed links to release notes and schedule for all the tools will be created on the EGI wiki, in the meanwhile a solution to get and publish this information in an (semi)automatic way will be investigated.

- 9) **Effort for specific adaptation of the Operations Portal:** The French NGI is devoting in effect 3 FTEs throughout out the duration of the project (but 68 PMs funded over 4 years) to cope with the requirements that lay on the Operations Portal’s team (those present from the beginning and the ones that the project naturally encountered on its way). Indeed, the team has met its requirements on the dashboard regionalization for operators across NGIs (ROD) and at the central level (COD), whereas in the mean time was also asked to deliver and delivered a security dashboard, latest feature request being availability/reliability reports. Also, since the beginning of the project it has been declared official EGI source for users metrics, and has been asked to develop transversal features that span way over pure operations (e.g VO dashboard).

Finally given the amount of features that are at its catalogue, the portal has a number of diverse and invisible tasks like support (see issue 5) that are unaccounted for, but still done because the team is devoted to this tool and because France has assigned that many FTEs throughout the whole project to this crucial tasks for the project. This is a minimum of permanent effort that is needed to provide a high quality tool, however the actual situation is by no way reflected in the PPT targets/consumption rate that has been assigned to it at the start of the project.

Mitigation: an analysis has to start/is planned within JRA1, SA1 and the project management to address this within PY3.

6 ROADMAP

The discussion of the roadmap for each of the JRA1 tools is beyond the scope of this deliverable and is already provided by a recent project milestone, the MS708 “Roadmap For The Maintenance And Development Of The Deployed Operational Tools” which details the roadmap for the next 18 months. However a short list of major expected developments in PY3 for each tool is provided in the following tables:

Operations Portal

	PQ9	PQ10	PQ11	PQ12/ PQ13
Major Upgrade of the regional package				
Refactoring of the Operations Dashboard				
First implementation of an Availability /Reliability module				
Mobile Version				

Table 4 – Operations Portal Roadmap for PY3

GGUS

	PQ9	PQ10	PQ11	PQ12/ PQ13	TBD
New Report Generator, 1st version					
High availability with auto-switching					
New Report Generator, production version					
High availability including Data Base Management System					
Specific work flows for CSIRT/Security					
Authentication through Shibboleth					

Table 5 – GGUS Roadmap for PY3

GOCDB

	PQ9	PQ10	PQ11	PQ12/ PQ13
Refining of Service Groups (Virtual sites)	█			
Redevelopment of XML output module	█			
GLUE2 compatibility	█			
Multiple endpoints per service	█			
Regionalisation reprioritisation			█	

Table 6 – GOCDB Roadmap for PY3

SAM

	PQ9	PQ10	PQ11	PQ12/ PQ13
Migration of grid middleware probes from SAM packages to the EMI packages	█			
MyEGI: Maintenance and stability improvements of the central monitoring service	█			
MyEGI: Improving the existing documentation (adding MyEGI user guide)	█			
POEM: improvements to the first version released in PQ8	█			
Messaging: client identification, check which instances are still using unauthenticated connections and setup credentials / request of clients to migrate to authenticated connections		█		
Messaging: Request to EGI OMB to approve authenticated only connections to PROD message broker network		█		
Messaging: Enforce authentication to all clients if OMB approves it			█	
Messaging: Enforce authorization rules if OMB approves it			█	

Table 7 – SAM Roadmap for PY3

Accounting Repository, including TJRA1.4, new resource types accounting

	PQ9	PQ10	PQ11	PQ12/ PQ13
Receiving Test Storage Records				
Improving cloud records handling				
Database migration				
Regional distribution of repository				

Table 8 – Accounting Repository Roadmap for PY3

Accounting Portal

	PQ9	PQ10	PQ11	PQ12/ PQ13
Maintenance of Accounting Portal 4.1 released on PQ8				
Contributed CPUs view				
Parallel Job accounting				

Table 9 – Accounting Portal Roadmap for PY3

Metrics Portal

	PQ9	PQ10	PQ11	PQ12/ PQ13
Version 3 for all the modules taking into account activity leader requirements				



Table 10 – Metrics Portal Roadmap for PY3



7 PARTICIPATION TO CONFERENCES AND PROJECT TASK FORCES

JRA1 and all its product teams actively attended the annual main project events, the Technical Forum and the Community Forum. During the Technical Forum in Lyon an accounting tutorial and an accounting workshop were organised and in various sections all the tools were discussed, in particular the SAM team gave a talk on the overall status and plans of the monitoring system and of the Availability Computation Engine (ACE). In addition a joint meeting with EMI and EGI on the technical details of the probes handover was organised.

JRA1, but in particular the GOCDB, actively attended the two “Towards an Integrated Information System” workshops, bringing requirements and presenting information system usage patterns of all the tools. GOCDB presented details about its architecture and plans for the future.

At the community forum an entire session was devoted to SAM, this session hosted a talk on the latest status of developments and a closed meeting for the training of the second level support team.

JRA1 participated to the FedCloud taskforce as already discussed in previous sections setting up the monitoring and accounting systems for the cloud federated resources.

JRA1 has been represented in all the project taskforces, groups and bodies that required its presence, i.e. new middleware integration taskforces, GGUS report generator group, the inter-project accounting taskforce etc. Finally a total of 5 OTAG meetings were held during PY2.

8 CONCLUSION

This document reported the second year activity of the project work package 7 (JRA1) within each of its tasks. WP7 deals with the maintenance and development of operational tools. JRA1 is composed of five tasks:

1. TJRA1.1 is the management task
2. TJRA1.2 for the maintenance and underline developments for all the tools
3. TJRA1.3 devoted to the development of regionalised tools
4. TJRA1.4 for the extension of the accounting system to encompass new resource types (other than CPU)
5. TJR1.5 for the development needed to the extension of the Operations Portal and its harmonization with other portal frameworks

The second year is, for JRA1, the only one with all its tasks active, in fact TJRA1.3 should have ended at the end of PY1 but was extend for another year because of underspending and to complete its activities, while TJRA1.4 started in PY2 as planned.

TJRA1.2 activities proceeded regularly during the year and many new features were added to all the tools, these are described in details in section 2 and among them we can mention:

- **GOCDB:**
 - a new finer grained roles schema
 - the implementation of service groups (also known as Virtual Sites)
 - the data scoping
- **SAM:**
 - the development of a new profile manager (POEM)
 - the integration of the probes for all the middleware types currently used in production
- **GGUS:**
 - a brand new report generator
- **Accounting system:**
 - a complete redesign for the accounting portal
 - a new data transport mechanism accompanied by a redesign of the accounting repository database
- **Metrics Portal:**
 - a complete redesign of the tool
- **Operations Portal:**
 - new modules added such as the Security Dashboard and the VO-oriented Dashboard.

For what concerns the integration of new (different from gLite) middleware types into the operational tools, this is now complete from the perspective of SAM and GOCDB for:

- UNICORE
- ARC
- GLOBUS
- Desktop Grids

TJRA1.3, devoted to regionalisation, was extended for one year at the end of PY1. During the second year developments for regionalisation proceeded slowly but the requirements were defined and analysed clearly and some progress in the development have been made in particular for what concerns GOCDB and its data scoping that allows to distinguish non EGI sites. Also the regional



package for the Operations Portal was improved during this year. However a full regionalisation for all the tools that were planned to be available for regional deployment is not available yet and will be carried on under the umbrella of TJRA1.2 in the coming years.

TJRA1.4 started in PY2 and, as first step, a survey on what is available in various projects and standardization bodies for the accounting of resource types different from the cpu usage was carried on. Focus was given to storage accounting and to definition of its usage record, to the accounting of the usage of virtual machines, and to parallel and mpi jobs. Some developments were also done for the accounting of cloud resources usage within the project FedCloud taskforce.

TJRA1.5 activities were reported in section 2 in the paragraph dedicated to the Operations Portal.

A series of main issues arose during the year and have been reported in section 5, we mention here the missing effort to face big requirements for the SAM team, the second level support for SAM that is done on a voluntary basis by too few people, the missing effort for regionalisation activities, the need for a inter-projects governing body for GGUS and the lack of a coordinated integration testbed for all the tools.

Section 5 also tried to provide a spent effort assessment for all the tools giving an estimate of the splitting between the effort needed for base maintenance and for new features developments. This is done in order to give guidelines for the funding of future activities that will deal with operational tools development. The maintenance effort tends to increase for almost all the tools and the rise is steeper for newer or completely redesigned tools such as the metrics portal (all the effort devoted to development in the first phase for these tools).

The PY3 roadmap for the tools is not described in details in this document, it was just briefly outlined in section 6 as it was the focus of a recent project milestone (MS708).

9 REFERENCES

R 1	JRA1 Main Wiki Page: https://wiki.egi.eu/wiki/WP7: Operational Tools
R 2	MS705 Operations Portal Roadmap: https://documents.egi.eu/document/525
R 3	Pakiti home page: http://pakiti.sourceforge.net/
R 4	The NAGIOS System Homepage: http://www.nagios.org/
R 5	EGI CSIRT Main Wiki Page: https://wiki.egi.eu/wiki/Csirt
R 6	Operations Portal Release Notes Browser: http://operations-portal.egi.eu/aboutportal/releaseNotesBrowser
R 7	MS410 EGI Helpdesk and the NGI support units: https://documents.egi.eu/document/522
R 8	GOCDB Regionalisation Plan: https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation
R 9	GOCDB Virtual Sites Design https://wiki.egi.eu/w/images/f/fe/VirtualSitesDesign.pdf
R 10	QoS CoS Grid Middleware: http://www.qoscogrid.org
R 11	IRODS Services: https://www.irods.org
R 12	GOCD NewRoles Plan: https://wiki.egi.eu/wiki/GOCDB/Release4/Development/NewRoles
R 13	MS 706 Operational Tools Accounting Work Plan: https://documents.egi.eu/document/531
R 14	Accounting Portal Roadmap: https://documents.egi.eu/document/517
R 15	Central SAM documentation https://tomtools.cern.ch/confluence/display/SAMDOC/Home
R 16	Gridmap https://openlab-mu-internal.web.cern.ch/.../EDS_GridMap_Poster.pdf
R 17	SAM Failover Configuration: https://tomtools.cern.ch/confluence/display/SAMDOC/Installing+SAM+Update-13#InstallingSAMUpdate-13-FailoverNagiosconfigurablehotstandbymode
R 18	Metrics Portal Roadmap https://documents.egi.eu/document/516

R 19	ActiveMQ message system homepage: http://activemq.apache.org/
R 20	GOCDB Service Types Query: https://gocdb.egi.eu/gocdbpi/public/?method=get_service_types
R 21	DMSU Main Wiki Page: https://wiki.egi.eu/wiki/TSA2.5_Deployed_Middleware_Support_Unit
R 22	EMI Project Homepage: http://www.eu-emi.eu/
R 23	IGE Project Homepage: http://www.ige-project.eu/
R 24	SAGA: www.ogf.org/documents/GFD.90.pdf
R 25	FedCloud Main Wiki Page: https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce
R 26	FedCloud Scenario5 Wiki Page: https://wiki.egi.eu/wiki/Fedcloud-tf:WorkGroups:Scenario5
R 27	EGI Technical Forum 2011: https://www.egi.eu/indico/sessionDisplay.py?sessionId=89&confId=452#20110922
R 28	Workshop on new types of Accounting at EGI Technical Forum 2011 https://www.egi.eu/indico/getFile.py/access?contribId=361&sessionId=89&resId=0&materialId=minutes&confId=452
R 29	UNICORE Integration TaskForce Wiki Page: https://wiki.egi.eu/wiki/UNICORE_integration_task_force
R 30	Inter Project Accounting TaskForce https://wiki.egi.eu/wiki/TCB:Accounting_Task_Force
R 31	Deployed SAM Instances https://wiki.egi.eu/wiki/SAM_Instances
R 32	GOCDB_Regionalisation_Plans https://wiki.egi.eu/wiki/GOCDB_Regionalisation_Plans
R 33	GOCDB Scoped Sites Plan: https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation#1.29_Data_Tagging.2FScoping
R 34	GOCDB Data Scoping Main Page: https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation/Data_Scoping
R 35	GOCDB Regional_Module_Technical_Documentation https://wiki.egi.eu/wiki/GOCDB/Regional_Module_Technical_Documentation
R 36	Operations Portal CVS https://cvs.in2p3.fr/operations-portal/regional/trunk/
R 37	D7.1- Annual Report on Operational Tool maintenance and development activity https://documents.egi.eu/document/372



R38

OGF-CPU Usage Record discussion wiki

<https://forge.gridforum.org/sf/wiki/do/viewPage/projects.ur-wg/wiki/ConsideredURFieldsCPU>

10 APPENDIX A1

Task	Partner	PY1		
		Worked PM Funded	Committed PM	Achieved PM %
TJRA1.2	10B-KIT-G	9,612	11,750	82%
TJRA1.2	12B-FCTSG/CSIC	5,600	3,000	187%
TJRA1.2	14A-CNRS	3,059	3,000	102%
TJRA1.2	16A-GRNET	0,914	3,000	30%
TJRA1.2	17-SRCE	3,290	3,000	110%
TJRA1.2	34A-STFC	5,269	6,000	88%
TJRA1.2	35-CERN	0,533	3,000	18%
TOTAL PY1		28,278	32,750	86%

Task	Partner	PQ5-PQ6-PQ7			TOTAL PY1+ PQ5/PQ6/PQ7		
		Worked PM Funded	Committed PM	Achieved PM %	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.2	10B-KIT-G	10,937	8,813	124%	20,549	20,563	100%
TJRA1.2	12A-FCTSG/CSIC	3,522	2,250	157%	9,122	5,250	174%
TJRA1.2	14A-CNRS	2,250	2,250	100%	5,310	5,250	101%
TJRA1.2	16A-GRNET	0,937	2,250	42%	1,851	5,250	35%
TJRA1.2	17-SRCE	2,634	2,250	117%	5,924	5,250	113%
TJRA1.2	34A-STFC	4,658	4,500	104%	9,927	10,500	95%
TJRA1.2	35-CERN	4,178	2,250	186%	4,710	5,250	90%
TOTAL PQ5/PQ6/PQ7		29,116	24,563	119%			
TOTAL PY1 + PQ5/PQ6/PQ7		57,393	57,313	100%	57,393	57,313	100%

Table A1 - Effort for TJRA1.2 in PY1 and PY2 excluding PQ8

		PY1						
Task	Partner	Worked PM Funded	Originally Committed PM	New Committed PM	Achieved PM %			
TJRA1.3	12B-FCTSG/CSIC	0,632	3,000	1,500	42%			
TJRA1.3	14A-CNRS	1,288	3,000	1,500	86%			
TJRA1.3	17-SRCE	3,567	3,000	1,500	238%			
TJRA1.3	34A-STFC	1,439	3,000	1,500	96%			
TJRA1.3	35-CERN	4,346	6,000	3,000	145%			
TOTAL		11,272	18,000	9,000	125%			
		PQ5-PQ6-PQ7				TOTAL PY1 + PQ5/PQ6/PQ7		
Task	Partner	Worked PM Funded	Originally Committed PM	Committed PM	Achieved PM %	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.3	12B-FCTSG/CSIC	1,030	0,000	1,125	92%	1,662	2,625	63%
TJRA1.3	14A-CNRS	2,178	0,000	1,125	194%	3,466	2,625	132%
TJRA1.3	17-SRCE	0,000	0,000	1,125	0%	3,567	2,625	136%
TJRA1.3	34A-STFC	0,146	0,000	1,125	13%	1,586	2,625	60%
TJRA1.3	35-CERN	0,000	0,000	2,250	0%	4,346	5,250	83%
TOTAL		3,354	0,000	6,750	50%			
TOTAL PY1 + PQ5/PQ6/PQ7		14,626	0,000	15,750	93%	14,626	15,75	93%

Table A2 - TJRA1.3 Effort for the period PY1 and PY2 excluding PQ8

		PQ5-PQ6-PQ7		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.4	10B-KIT-G	5,7	4,5	127%
TJRA1.4	12A-CSIC	0,130	4,500	3%
TJRA1.4	21A-INFN	0,905	6,500	14%
TJRA1.4	34A-STFC	5,885	6,750	87%
TOTAL PQ5/PQ6/PQ7		12,636	22,250	57%

Table A3 - TJRA1.4 effort for PY2 excluding PQ8. During PQ1 the task was not active.



Task	Partner	PY1		
		Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	13,488	17,667	76%
PQ5-PQ6-PQ7				
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	12,050	13,250	91%
TOTAL PY1+ PQ5/PQ6/PQ7		25,537	30,917	83%

Table A4 - TJRA1.5 effort for PY1 and PY2 excluding PQ8