EGI-InSPIRE

D7.3 ANNUAL REPORT ON OPERATIONAL TOOL MAINTENANCE AND DEVELOPMENT ACTIVITY

EU DELIVERABLE: D7.3

Document identifier:	D7.3-1562-FINAL
Date:	**01/05/2013**
Activity:	**JRA1**
Lead Partner:	**INFN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/1562

Abstract

The activity that has taken place during the 3rd year of the EGI-InSPIRE's work package 7 is reported in this document together with an assessment of the effort spent on all of the tools. The roadmap for the next year of each tool is also reviewed but not elaborated in details since it was the focus of a recent project milestone.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Diego Scardaci	INFN	18/03/2013
Reviewed by	John Kennedy Mark Santcroos Maarten Litmaath	RZG AMC CERN	17/4/2013 24/4/2013 23/4/2013
Approved by	**AMB & PMB**		**30/4/2013**

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	18/03/13	First draft	D. Scardaci
2	02/04/13	Internal Review	D. Scardaci
3	24/04/13	External Review	D. Scardaci

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 third year activity in work package 7 (JRA1) of the EGI-InSPIRE project, which deals with the maintenance and development of operational tools.

The routine maintenance and development of the operation tools proceeded regularly during the year and many new features were added to all the operational tools. As a result, EGI is able to support within its monitoring and service registration capabilities, services from:

- UNICORE
- ARC
- GLOBUS
- Desktop Grids
- QosCosGrid

In addition, operational tools and other EGI services, as the Applications Database, the Training Marketplace and the Customer Relation Manager are now monitored by the operational tools.

In PY3 regionalisation activities have continued following an analysis of the benefits to having a regionalised version of some of the tools taking into account the real NGIs needs. In some cases it has been found that creating regional views on the central instances is an acceptable alternative to providing the effort needed to develop and maintain a regionalised version of the tool. As a result, development on the Operations Portal and GOCDB regional versions has stopped in favour of regional views on the central instance. Developments on the regional accounting solution continued during PY3 and the regionalised release of the Accounting Repository and Portal will be deployed in production, respectively, in April and May 2013.

After spending the previous year focused on investigation, work started over the last year on implementing accounting for different resource types. This initial development work means that within the next year Cloud, CPU and Parallel Jobs resources types could be accounted for in an updated EGI Accounting Repository, and initial design has taken place to provide accounting support for applications, storage and virtualisations.

The dedicated development work relating to new capabilities in the Operations Portal has come to an end this year and routine maintenance activities will now continue.

The document provides an analysis of the overall effort consumption for each operational tool and the distribution made between maintenance and development work. The maintenance effort has grown for almost all tools during the project's lifetime. This reflects the decrease in the number of new features being developed once most user requirements have already been fulfilled and the increase of the effort needed to refine the already delivered capabilities (e.g. bug fixing). Exceptions to this trend are the GOCDB and the Accounting Portal due to the development of GOCDB v5 and the integration of the new resources types in the Accounting Portal which has increased the development load related to maintenance. This shift may indicate that the operational tools are now entering a phase of consolidation rather than new developments.

The roadmap for the operational tools in PY4 is not described in detail in this document as it has been the focus of a recent project milestone (MS710).

TABLE OF CONTENTS

1 INTRODUCTION

This document reports the third project year activity of the 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 the 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. TJRA1.5 for the extension of the Operations Portal and its harmonization with other portal frameworks

TJRA1.5 completes its activities at the end of this year, while TJRA1.3 ended in PY2 after one year extension.

In this deliverable there is also an effort assessment for each tool, in attempt to distinguish the effort spent in maintaining the tools and the effort used to develop new features. This is done in order to provide guidelines for funding of future activities/projects that will maintain the operational tools. The PY4 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 (MS710). A series of issues have arisen during the year and are also reported.

The document is organised into the following main sections:

- section 2 refers to task TJRA1.2 and describes the maintenance activity and the development needed to address the main user requirements for all the tools. Regionalisation activities are also reported in this section;
- section 3 provides details about TJRA1.4 and the accounting for new resource types;
- section 4 reports the spent effort assessment and the main issues encountered during the year;
- section 5 is a quick overview of the next year roadmap for all the tools;
- section 6 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 7.

2 DEVELOPMENT AND MAINTENANCE

This section gives a summary of the development and maintenance activity performed for each of the tools in the JRA1 during PY3. The JRA1 contribution to the implementation of the EGI Infrastructure Platform and Collaboration Platform described in D4.6 has been highlighted[1]. A very short description of each tool is also provided at the beginning of each section together with references to outline tool architecture, documentation, deployment scenarios and future developments.

The general entry point to get information about the JRA1 activity and its tools is hosted by the EGI wiki at [R 1].

2.1 Operations Portal

The Operations Portal provides information to various actors (NGI Operations Centers, 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 that can host regional views. A detailed description of the whole system and of its modules can be found in [R 2].

The development of the Operations Portal during the entire PY3 focused on the following components:

1. **Monitoring of unsupported middleware version**: security and operation dashboards are now able to expose the alarms raised when an obsolete version of a middleware component is detected;
2. **Dashboards**: an important refactoring has been initiated to improve their performance and usability and a new VO Operations dashboard has been designed;
3. **Availabilities/Reliabilities for Top-BDII and sites**: these modules have been designed to compute and expose the availabilities and reliabilities of the Top-BDII of a NGI and sites;
4. **VO Availabilities/Reliabilities**: the Operations Portal team is developing a new module able to compute VO specific availability/reliability reports;

The security and the operations dashboards have been modified to expose the **alarms raised in cases of unsupported middleware version detection** to support the EGI Operations team to identify obsolete versions in the deployed middleware components.

The **dashboards refactoring** has been planned to provide a new portal look and feel with a homogenization of the display and an improvement on efficiency, reactivity and visibility. A first prototype has been completed in December and delivered to ROD Operators for a testing phase. The production version will be released by June 2013 and will be properly updated to fulfil the ROD Operators requirements collected during the testing.

The Operations Portal is now able to provide **availability and reliability information for Top-BDII and sites**. The lists of NGI Top-BDIIs and sites are respectively acquired parsing the EGI SA1 wiki [R 3] and from the GOC DB programmatic interface. Then, the MyEGI interface is queried to obtain the input data to compute A/R statistics with the designed algorithm [R 4]. A/R data are stored in the Operations Portal database and exposed by providing tables, charts and possibility to export data. Figure 1 shows the process to compute Top-BDII A/R in graphical format.

[1] See https://documents.egi.eu/secure/ShowDocument?docid=1309

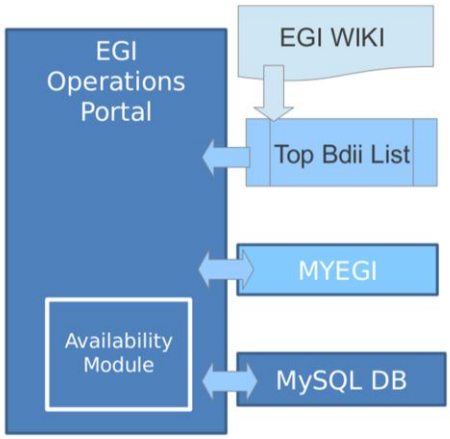

Figure 1 - Process to compute Top-BDII A/R. The EGI WIKI is replaced with the GOCDB to compute sites A/R.

The JRA1 management, in collaboration with all the product teams involved, analyzed the requirement identified at the last project review to provide more user oriented availability and reliability (A/R) information through **monthly VO A/R statistics** [R 5] that would reflect more the experience of a VO. As a result of this study a design document [R 6] has been written defining the method to compute VO A/R and the actions that should be performed by the PTs involved to fulfil the requirement.
A method, was devised assuming the following guidelines:

- VO specific reports are generated monthly, using the OPS VO;
- Only reports for "high activity" VOs are generated. With high activity we can follow the definition of the accounting portal, i.e. high activity week if
  - CPU time consumed over a week > 1 year;
  - CPU time consumed over a month > 4 year;
- The list of service end-points supporting the high activity VOs needs to be extracted from Top-BDII.

The workflow of the agreed method is described in the following and in Figure 2:

- The Accounting Portal builds the High Activity VO list;
- The Operations Portal queries the Accounting Portal to get the High Activity VO list;
- For each VO of the High Activity VO list, the Operations Portal queries the BDII and extracts a list of services for it;
- The Operations Portal groups the services by type and filters the list of services with the ones computed by SAM;
- For each service in the list, the Operations Portal queries SAM to retrieve raw A/R data per services on a daily basis;
- The Operations Portal computes the metrics with the dedicated algorithm.

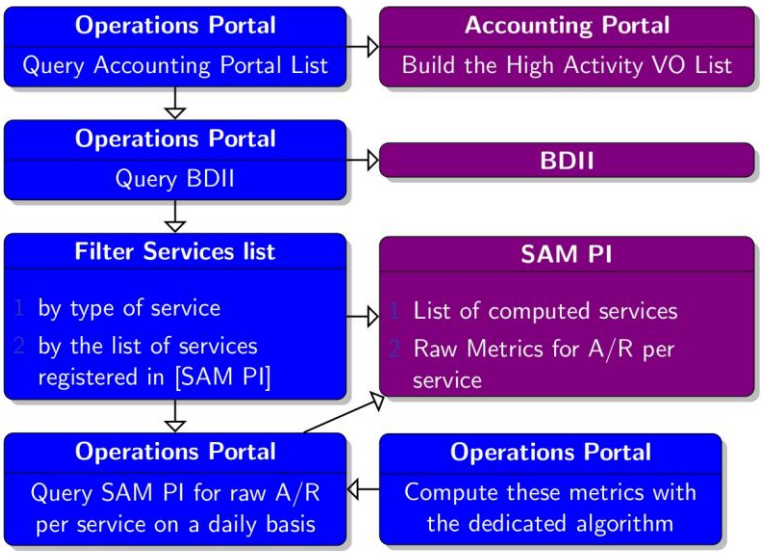

Figure 2 - VO Availability/Reliability – Method workflow

The algorithm chosen to compute these metrics can be summarized in the following points:

1. For every service group[2] calculate the availability/reliability value by ORing the A/R values of the services contained in the group;
    a. The requested time granularity of this computation has to be assessed, in principle a daily granularity is acceptable (hourly would be optimal);
    b. For the first implementation a plain OR is what is requested, for future implementations a 'threshold conditioned OR' would be more suitable to implement VO SLAs;
2. Perform a plain AND of the groups availability and reliability to calculate the overall availability/reliability of the infrastructure
    a. This will require to be computed with the same time granularity used in the OR step

This algorithm does not provide information about the perceived status of the infrastructure from an individual user's point of view, but does provide a re-aggregation of the data from a VO perspective. In the future more complex algorithms could be adopted that take into consideration averages and weights for each kind of resource (e.g. a VO supported by 100 CEs may require at least 30 CEs available to consider the e-Infrastructure usable).

The new Operations Portal module to compute A/R VO reports will be released by April 2013. A prototype is already complete (see Figure 3).

[2] A service group is a set of the same resource type (e.g. CE, SE) supporting a VO.

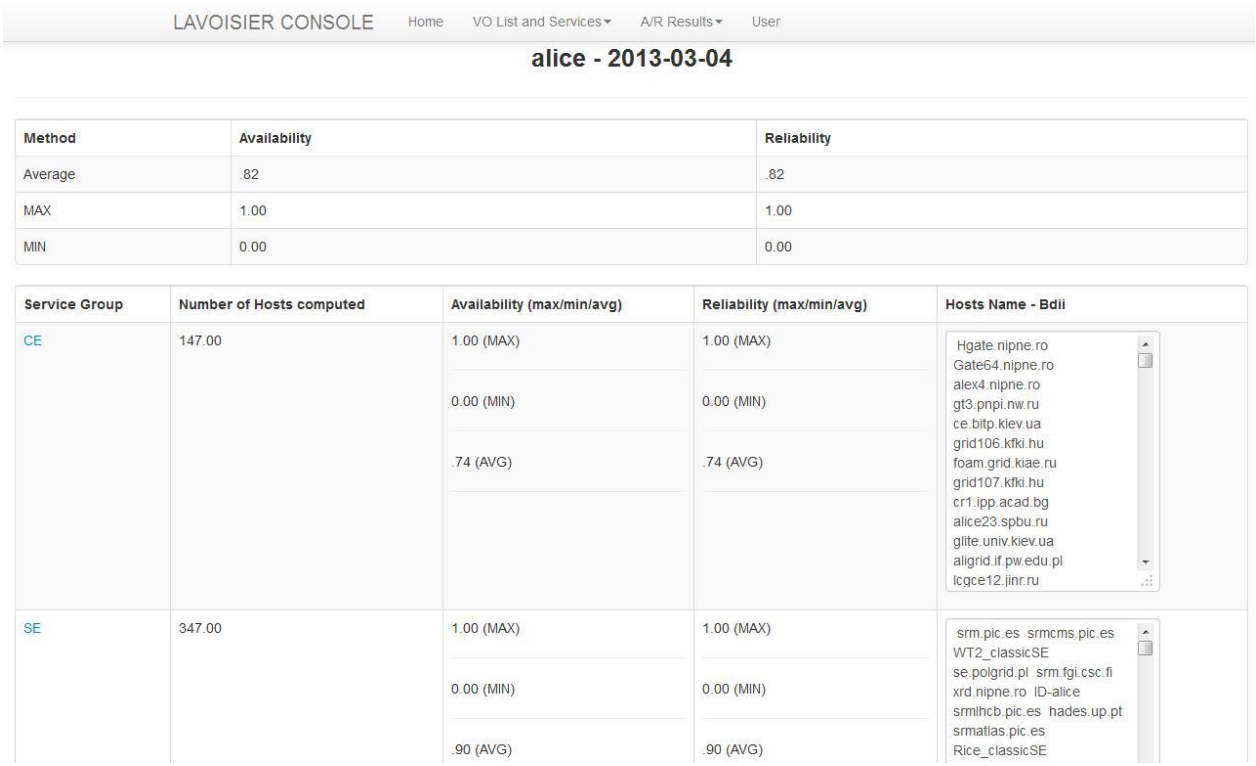

Figure 3 - Operations Portal module to compute VO A/R – Prototype.

The Operations Portal team proposed, during OTAG-13, to drop the **Regional Package** of the Operations Portal taking into account the following considerations:

- the number of functionalities and modules requested and developed for the tool are increasing, dropping the regionalised version could help to maintain the high-level service provided by the central instance;
- three of the four NGIs that deployed the regional package asserted (via mail survey or directly at the OTAG) that the added value of the regionalised version of the Operation Portal is not that big and their needs could be addressed also by the regional views hosted in the central instance.

During the last months of PY3 those four NGIs tried the new regional views on the central instance and verified that their needs are well satisfied by them. These four NGIs have came back to the Operations Portal central instance and, as consequence, the regionalised version of the Operations Portal has been dropped.

During PY3 a total of 3 releases were performed by the Operations Portal team, detailed release notes are available at [R 7].

2.2 EGI Helpdesk (GGUS)

GGUS is the EGI helpdesk [R 8] and the single 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 third year:

1. The **Report Generator** has been redesigned from scratch. This new version provides the ability to create all metrics reports on the fly;
2. The **High Availability (HA) solution** has been implemented for the Web Front-ends and the AR Server (BMC Remedy Action Request System). Moreover, HA is now available for the on call duty service and the Intrusion Prevention system;
3. A new **GGUS advisory board** has been established during the EGI Technical Forum in Prague [R 9]. It is composed of representative from user communities, NGIs, EGI, technology providers. The advisory board meets once per month directly after the monthly GGUS releases. The new requirements will be discussed there before reaching the OTAG;
4. Several **minor achievements** have been completed such as the adaptations for the interface to the CERN ServiceNow ticketing system, the interface with the NGI France ticketing system, the integration in the GGUS helpdesk system of new support units for some VOs and a new GGUS mail infrastructure.

The definition of the new interfaces to PRACE/MAPPER [R 10], DANTE [R 11] and to the NGI Ibergrid is ongoing, as is the implementation of specific work flows for CSIRT/Security [R 12].

In addition to the already available authentication methods, X509 certificates and login/password, new authentication methods to access GGUS services are under examination. Indeed NGIs expressed interest in federated identity management technologies such as Shibboleth.

The **xGUS helpdesk** [R 13] covers the regionalisation needs of the NGIs that do not have a custom solution for their own local ticketing system. The xGUS helpdesk template has been developed for NGIs and user communities who want to build up their own user support infrastructure. It contains all basic helpdesk functionalities including user administration and certificate access and allows the creation of customisable views.

The GGUS product team published a total of 10 releases during PY3[3].

2.3 Grid Configuration Database (GOCDB)

The GOCDB is the main source of topological information for EGI. 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 PY3 effort focused on the following main activity streams:

1. In September 2012 the product team released the **GOCDB v4.4**. This release harmonized the separate read-only and read/write instances into a single portal and addressed many small RT requirements for GUI enhancements. See the GOCDB change log for more details [R 14].
2. A **GOCDB failover instance** has been installed at the Fraunhofer Institute [R 15] to increase the tool availability.
3. A **separate GOCDB instance** has also been **deployed by the EUDAT project** [R 16], and has been tailored for use by means provided by the GOCDB abstractions/extension-points.

[3] Since many years GGUS releases usually are done on the last Wednesday of a month. The high frequency of releases has the advantage to be able react quickly to the many requests coming from various groups and communities.

4. The last two PY3 quarters have largely focussed on the **design and development of GOCDB v5**. This new GOCDB version is based on a new data layer able to use different RDBMS platforms (e.g. MySQL, Postgres, Oracle) satisfying the DoW requirement [R 17] to support non-Oracle deployment. It will simplify future developments making GOCDB more attractive for adoption by other projects. GOCDB v5 is aiming for its first production release in May 2013.
5. Continued engagement with the GLUE2 working group to help to finalize the GLUE2 XML rendering document. This has since been **submitted to OGF for public comment.**

Several other activities are ongoing or planned for PY4 to **increase the appeal of GOCDB for use by other projects**. The GOCDB scoping will be extended to introduce **multiple, non-exclusive scope tags** to enable hosting multiple projects within a single GOCDB instance[4]. The **extensibility mechanism** will allow the core GOCDB entities (NGIs, Sites, Services, ServiceGroups, Endpoints) to define an optional set of custom key-value pairs (i.e. custom property bags). This could potentially be used for VOs and other custom/domain-specific fields. This activity is not foreseen in the original plan for JRA1 and has been funded by EGI-InSPIRE in the context of the call for mini-projects of January 2013 [R 18] and will not take place within SA4. The post v5 releases will include also **GLUE2 support** and new PI method(s) to programmatically input downtime information.

Also the **regional package of the GOCDB** has been dropped in agreement with the NGIs taking in consideration that:

- The usage of local scoping in the central GOCDB instance addressed the majority of NGI requirements for adding non-EGI sites and services within the central GOCDB;
- The same NGIs had a limited effort/interest to host and manage a separate regional GOCDB (revealed after a more detailed NGI poll);
- Technical complexities due to the existence of a distributed GOCDB hierarchy (regional instance vs. central instance) would have considerably increased the effort needed to develop and maintain the GOCDB. For example the scope for data inconsistencies between central and regional instances would be significantly increased and publishing/latency issues would have complicated matters further.

The GOCDB product team published two major releases during PY3 whilst also continuously deploying a number of functional enhancements throughout the period.

2.4 Accounting Repository

The EGI accounting repository [R 19] stores information relating the usage of resources within EGI's production infrastructure. It receives data on individual jobs and summaries of collections of job records from information providers, sites and other infrastructures, and exports accounting information to consumers of usage records, for example the accounting portal for visualization. Control of which sites publish to the APEL repository is defined by information pulled from GOCDB. Monitoring of whether sites actually publish is carried out centrally by APEL and the results are published to SAM.

The most substantial result obtained during the third year of the project has been the bringing into service of the **new SSM (Secure Stomp Messenger)** [R 20] **based APEL Accounting Repository** in June 2012. Sites of some big EGI partners, as CERN and NIKHEF, now send Job Records to this

[4] The current scoping implementation provides only 'Local' and 'EGI' scopes which are mutually exclusive.

repository and external clients, as OSG/Gratia, INFN/DGAS, NDGF/SGAS and Switzerland/SGAS, have migrated to sending Summary Records (for 97 sites) to the new APEL Accounting Repository. New SSM clients have been developed and tested (or are in testing phase) with the Accounting Repository for IGE/GridSafe [R 21], QoSCoSGrid/MAPPER [R 22], EDGI [R 23] and UNICORE [R 24]. The records collected by the new APEL Accounting Repository are integrated with the summaries from the old Accounting Repository and retrieved by the Accounting Portal.

In March 2013 EMI-3 was released. It includes the new EMI-APEL client, completely rewritten, which will use SSM v2 for communication between clients and the APEL Accounting Repository and includes support for local jobs and MPI accounting. When the sites are updated to EMI-3, they will automatically send their records to the new SSM based APEL Accounting Repository which uses the EGI Message Brokers to receive data. The previous version of the APEL Accounting Repository will be retained while support for the old EMI versions continues unless all sites migrate earlier. During this period the two accounting repositories will run in parallel and the Accounting Portal will retrieve data merged from both repositories.

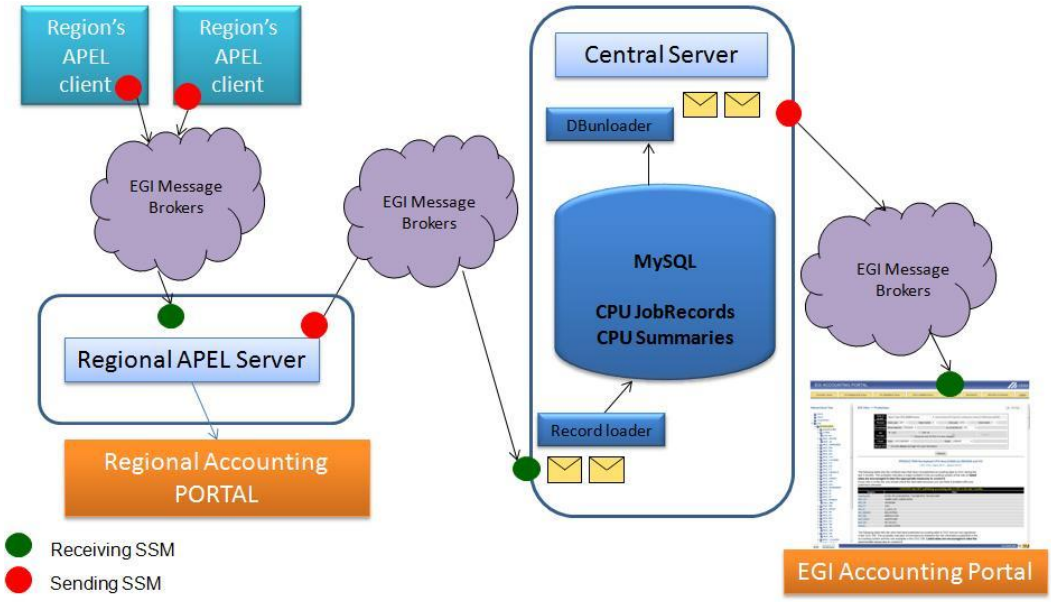

Figure 4 - Regional Accounting Repository integrated in the whole accounting ecosystem.

The **Regional Accounting Repository** will be released into production by April 2013. It comprises:

- MySQL database;
- apel-lib, apel-server and apel-ssm packages;
- Packages are available on github and documentation for installing a test setup is available.

The communication between clients and servers is done via the EGI Message Broker network using the APEL SSM package like for the new central repository. The New EMI 3 APEL Client can be configured to send Summary Records to either the central APEL server or a regional APEL server. **Error! Reference source not found.** depicts how the Regional Accounting Repository will be integrated in the whole accounting ecosystem.

2.5 Accounting Portal

The data recorded in the Accounting Repository (section 2.4) are processed, summarized and pulled to 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 understand resource utilisation.

During PY3 the Accounting Portal team **improved the product with a considerable code refactoring and several enhancements and optimizations mostly driven by user requirements**. The most important improvements are the following:

- Extension and maintenance of the VO Manager views;
- User data views now show all user entries;
- Graphs made bigger and easier to read;
- Support of RFC2254 DNs;
- InterNGI usage reports;
- PDA & Mobile support;

Two other important activities will be completed by the end of PY3. A **preliminary support for the provisioning of Cloud Accounting** and the **XML endpoints generalization and improvement**, that would be documented and made public, will be deployed in production in April 2013.

The current production version of the Accounting Portal is **v4.2 Fomalhaut**. New production releases will be available every 6 months.

The **Regional Accounting Portal** will be deployed in production by May 2013. It will be distributed as a preconfigured, ready-for-use VM and the updates will be published via a code repository, improving solution time and enabling local branches with regional specific code. The NGI staff will be required to do small administration tasks only. The Accounting Portal team will assist NGIs and will have a copy of their installation in hand to debug and propagate fixes. The Regional Accounting Portal will display local sites and federations, so some views (e.g. OSG) and reports will be removed. Each VM will be tested against the corresponding regional APEL server.

2.6 Service Availability Monitor (SAM)

The Service Availability Monitor (SAM) [R 25] is the system that is used to monitor EGI resources within the production infrastructure. It consists of the following components:

- Probes: for the individual service component being monitored in the test;
- Submission framework: a test execution framework (based on the Nagios open source monitoring framework) and the Nagios Configuration Generator (NCG);
- Storage layer: 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.

The main activity streamlines on which the development focused during PY3 are:

- **POEM** component has been included and **fully integrated** in SAM Update-17 (Release notes available at [R 26])

- **MyEGI** has been **reviewed** and improved as part of SAM Update-19 (Release notes available at [R 27]). It is currently providing the following views: Metric Status, Availability and Reliability, Treemap, and Topology description.
- **SAM instance for monitoring operational tools** (OPS-MONITOR) has been deployed as part of SAM Update-20 (Release notes available at [R 28])

A Central OPS-Monitor instance has been deployed for monitoring operational tools. It publishes results to the central monitoring service where the status of the OPS services is computed.

An important activity started in PY3 is the **integration of EMI probes in SAM** [R 29]. This currently involves rearranging and validating meta-packages to support dependencies provided by EMI. This will also involve dependency changes in other SAM components to reflect the new arrangement as well as code adaptations that may be required for this transition. A working group, the EGI SAM probes WG [R 30], composed of experts from NGIs, EMI and EGI, was created to revise EMI probes before they are integrated into SAM framework and to evaluate probe and monitoring-related improvements. The activity of this WG will continue after the end of the EMI project. The integration of some EMI probes will be completed in September 2013.

SAM fully supports the EGI regionalisation plan from the end of PY1 and **all the NGIs are running their local SAM regional instance** [R 31]. All NGI instances are configured to use ATP as topology source. All DoW SAM regionalisation requirements for TJRA1.3 are now addressed, the new version of the MyEGI portal is able to provide EGI-specific NGI views and is also deployed on the regional instances. 'GridMap style' TreeMap views have been added, showing both regional and global views.

A total of two SAM updates were released to production during PY3, one in PQ10 and one in PQ11. All updates went through the full EGI Software Provisioning process [R 32] as was decided at the beginning of the project.

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 automatically collects the relevant data that are published by the middleware or by other operations tools (e.g. GGUS ticket statistics) via a series of connectors. In addition, some metrics are inserted manually by the activity leaders or by the NGI managers.

The main Metrics Portal developments performed during PY3 were:

- Per country metrics for NGIs that requested this feature;
- Heavy query optimization;
- Added XLS output support;
- Aggregated metrics (sum of all NGI predicted metrics plus entered metrics);
- Metrics accessible depending on several variables;
- Internal documentation and re-factorization;

The Metrics Portal has been used for the last two years to gather metrics from the project tasks. It has been updated according to changes in the structure and scope of the project and its tasks and activities.

In PQ9 new metrics for the SA2, NA2, SA1 and NA3 tasks have been added. The concept of depreciable metrics or activities has also been introduced.

2.8 Broker network configuration and infrastructure

The JRA1 activity is responsible for providing support for the configuration and operation of the message broker network of the production infrastructure based on the ActiveMQ system [R 33]. The main achievements for this activity during PY3 are:

- Implementation of the credential synchronization system, which is responsible for keeping user and group records synchronized between the brokers of the same broker network;
- Development of a test suite in order to test the message brokers network prior to applying software updates on production message broker network;
- Enabled the logging of unauthenticated connections (IPs) to the production broker network;
- An analysis with SAM team to implement authentication, preferably based on usage of X.509 credential, has started in December 2012;
- Redesign of monitoring tools (i.e. Nagios probes for SAM) has started in January 2013;
- An analysis to develop a failover capability with respect to the delivery of results from SAM probes has started: when a broker endpoint is not functional, the probe should be able to deliver its results to another broker endpoint within the network. This feature will be ready in December 2013.

2.9 Integration of new middleware types

The 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.

The number of service types defined in the GOCDB continued to increase and now we have 94 service types registered (66 at the end of PY2) from the following middleware stacks (a full list of services types can be obtained from the GOCDB programmatic interface [R 34]): gLite, UNICORE, Globus, iRODS [R 35], ARC [R 36], QosCosGrid, **BES** [R 37], **Cloud**, **Torque** [R 38], **Squid** [R 39], **XRootD** [R 40] (in bold the middleware stacks added in PY3). In addition to the operational tools service types, already present in GOCDB at the end of PY2, the last year has seen the registration of other EGI tools like the EGI Applications Database, the Training Marketplace, the Customer Relation Manager, etc. Moreover, new services types have been registered to identify portals and frameworks to access grid resources.

All new service type requests need to be assessed by EGI via a lightweight review process (by OMB and OTAG) so that only suitable types are added to GOCDB and to prevent duplication. The procedure to request a new service types is described in the GOCDB documentation [R 41].

The SAM monitoring framework currently is able to monitor services from the following middleware stacks: gLite, UNICORE, Globus, ARC, QosCosGrid and Desktop Grids. New probe integration for QosCosGrid QCG/MAPPER, UNICORE Job and unicore6.StorageFactory has been completed during PY3. EGI services and tools are considered as service types and probes are integrated in SAM in order to check their availability.

The 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 middleware 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 [R 42], EMI [R 43], IGE, SAGA[R 44], QosCosGrid and EGI

Federated Cloud [R 45]. The procedure to create a new support unit is described in the GGUS documentation [R 46].

Concerning the tools to operate the federated cloud infrastructure, JRA1 representatives continued to attend and actively participate in the EGI Federated Cloud task force activities. JRA1 led two working groups (also called Scenarios), one for monitoring (Scenario 5 [R 47]) and one (Scenario 4 [R 48]) for implementing an accounting system of the cloud infrastructure.

3 ACCOUNTING FOR DIFFERENT RESOURCE TYPES

Supported by the EGI-InSPIRE project, several new resource types have been integrated into EGI's production infrastructure and JRA1 has worked to include them in EGI's accounting system.

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 [R 49].

After a first year 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, TJRA1.4 mainly spent PY3 developing the solutions defined at the end of PY2.

Within this project year **the following resources types could be accounted in the new EGI Accounting Repository based on SSM**:

- Cloud;
- CPU;
- Parallel Jobs;

and significant steps have been performed to gather accounting data for:

- Applications;
- Storage;
- Virtualisations;

The JRA1 accounting product team is also participating in the **TCB's Inter Projects Accounting Task Force** [R 50] to get requirements about the integration of new middleware stacks (beyond ARC, UNICORE, gLite and Globus).

Another TJRA1.4 target is to evolve the accounting system to support an economic model that could contribute to the sustainability of EGI. However, the relevance of the inclusion of the new resources types in the accounting system for the EGI users and, in particular, the steadily increasing importance of the Cloud world, gained day after day during the project lifetime, pushed JRA1 to mainly devote TJRA1.4 effort to these activities. In this way we could fulfill the captured user requirements. In PY4 an analysis will be performed to understand how TJRA1.4 could help EGI.eu in the implementation of the Pay-for-Use model presented in the Evolving EGI Workshop (28-30 January 2013) [R 51]. Moreover, TJRA1.4 has joined the EGI Pay-for-Use Pilot Group [R 52].

3.1 Clouds

The JRA1 accounting product team continued participating actively in the work of the project's Federated Cloud taskforce and led the so-called "Scenario4- Accounting" activities with the mandate to include into the accounting system also the records coming from the Federated Cloud infrastructure.

During PY3, a new version of the Cloud Accounting Usage Record has been defined according to the results of the taskforce analysis. **The new Accounting Repository based on SSM has been deployed and successfully used to test the collection of the cloud accounting records**. Now, several cloud clients are sending their accounting data to the new Accounting Repository. These clients come from OpenNebula [R 53] and OpenStack [R 54] systems, while an investigation to include other systems, as WNoDeS [R 55] and Okeanos [R 56], will be done next year.

The Accounting Portal team is developing the visualisation of the cloud accounting data in the portal. In April 2013, the Accounting Repository will be able to send Cloud accounting summaries to the Accounting Portal and, at the same time, a new version of the Accounting Portal will be released

including a preliminary support of visualisation of Cloud Accounting. In October 2013, the Accounting Portal will include an enhancement of this support.

3.2 Parallel Jobs

After the definition of the parallel jobs accounting usage record within standardization bodies (i.e. OGF UR-WG [R 57]) last year, with the active participation of JRA1, the team started the development in 2012. **The parallel jobs data** has been added to the CAR (Compute Accounting Record) and, **after the release of EMI-3** that includes the new EMI-APEL client, such data **can be stored in the Accounting Repository**.

The Accounting Portal team is developing the visualisation of the parallel jobs accounting data in the portal that will be completed in October 2013.

3.3 Storage

The JRA1 product team is also continuing to participate in the definition of the storage accounting record within standardization bodies (i.e. OGF UR-WG). Meanwhile **an interim standard (Storage Accounting Record – StAR) has been defined** in EMI [R 58] reflecting practical, financial and legal requirements concerning storage location, usage and space and data flow. The definition might be the base for a standardized schema or an extension of an existing record like the OGF UR. During PY3 the JRA1 team **integrated the new record in the Accounting Repository and is now able to receive it through SSM from dCache, StoRM and DPM** storage types. An analysis to support StoRM storage is running, as it currently does not include support for storage accounting however, by exploiting the BDII information that it publishes, an approximate sensor could be developed. The Accounting Repository will be able to send storage accounting summaries to the Accounting Portal in September 2013.

The feature to visualise the storage accounting data is under definition and will be added in the Accounting Portal in October 2013.

3.4 Others

TJRA1.4 team started the **definition of an Application Accounting usage Record (AAR)** in XML format and developed a first prototype for Application Accounting. The AAR definition will be completed in April 2013. The integration in the Accounting Repository and Portal will be done in PY4.

4 EFFORT ASSESSMENT AND MAJOR ISSUES

4.1 Effort Overview

Task composition of the activity was discussed in details in D7.1 [R 59] 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**: Specific development activity 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 some EGI sustainability scenarios
- **TJRA1.5: Integrated Operations Portal**: Dedicated to the development needed for a restructuring of the Operations Portal, its evolution and harmonization with other portal frameworks

The JRA1 tasks follow the time sequencing described in Figure 5.

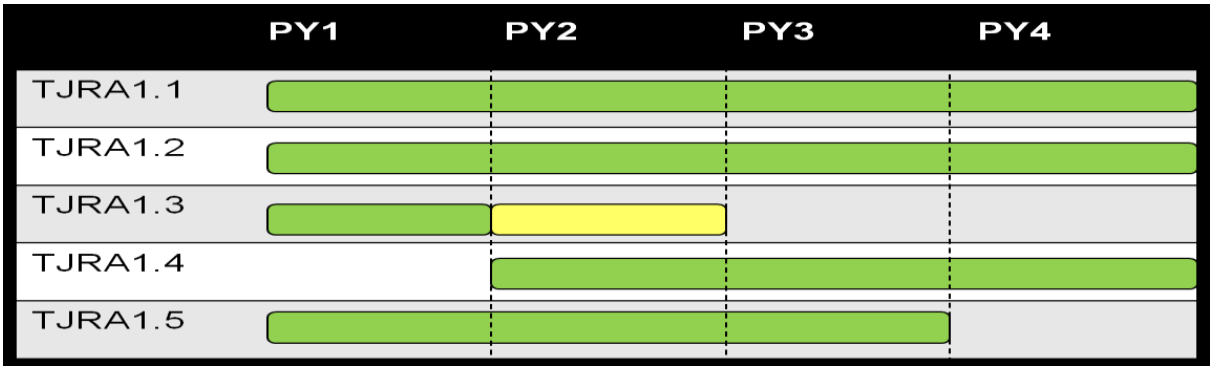

Figure 5 - JRA1 tasks time sequencing. TJRA1.3 was extended for another year at the end of PY1.

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 Appendix A. They show the data for the first two years of the project, for the first three quarters of the third year and for the sum of the three periods when possible (some tasks may have been inactive in one of the timeslots).

Task by task the following considerations apply:

TJRA1.2

The total TJRA1.2 effort consumption is in line with the committed effort. The main deviations highlighted at the end of PY2 are even now present but they have been partly mitigated during PY3. GRNET under-reporting has been reduced during PY3. FCTSG/CSIC over-reporting will be naturally mitigated in PY4 when FCTSG/CSIC will focus its effort in TJRA1.4 activities as planned in the roadmap defined in MS710 and summarized in Section 5 of this document.

TJRA1.3

The task ended at the end of PY2. The remaining activities have been moved to TJRA1.2.

TJRA1.4

TJRA1.4 still shows a rather significant under-reporting although it has been partly mitigated during PY3. This under-reporting can be considered fairly natural taking into account the type of activities planned for this task, activities that requires a deep study and a careful requirements collection before starting with the development phase. FCTSG/CSIC underreporting has already been mitigated during PY3 and will be strongly reduced during PY4 when they focus their effort on TJRA1.4 activities as explained above. An activities roadmap for INFN is going to be defined by the end of PY3 to considerably reduce its underreporting. The actions described above should largely compensate for the task underreporting before the end of the project.

TJRA1.5

TJRA1.5 completes its activities at the end of PY3 and shows no large deviations from the committed effort. There could be a slight underreporting due to a developer leaving the Operations Portal team in December and for whom a substitute has not been identified yet. After the development of several major features during the first three years of the project described in the previous JRA1 deliverables and Section 2.1, the Operations Portal team will be devoted mainly to maintenance activities, in the context of TJRA1.2, during PY4.

4.2 Development vs Maintenance Effort

The evaluation, tool by tool, of how the used effort is split between basic maintenance (bug fixes, interaction with other tools) and development for new features or code restructuring continued during PY3. This estimation allows us 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. The evolution of both effort numbers for the entire duration of the project could at least show a trend of the maintenance cost of the tools. Cost in this context is used in terms of effort and not of money, as translating effort to money is outside the scope of this deliverable.

The split between maintenance and development has been self-assessed by each of the product teams for each of the tools. This assessment for the first three years of the project, tool by tool, together with a forecast for PY4 that will need to be validated during the next project year are reported in Table 1, Table 2 and Figure 6. Table 1 shows the absolute numbers, tool by tool, of the total project effort used (in PY1, PY2 and PY3), while Table 2 reports the splitting between development and maintenance.

Throughout the project the focus work has changed from the initial development phase to the latter deployment and support phase and, for this reason, the maintenance effort grows for almost all tools during the project lifetime. This reflects the decrease of the number of new features to be developed once several users' requirements have already been fulfilled and the increase of the effort needed to refine the already delivered capabilities (e.g. bug fixing). The maintenance effort will start to decrease when the new functionalities become more stable and the features set quite frozen.

Exceptions to this trend are the GOCDB and Accounting Portal.

The GOCDB roadmap foresees the release of v5 in the first months of PY4. This activity includes the development of a new data layer able to use different RDBMS platforms (e.g. MySQL, Postgres, Oracle) and thus requires a considerable development effort for the next year.

The Accounting Portal team defined a tough roadmap to include the provisioning of new resource types that will be available in the Accounting Repository during PY4, such as Cloud, Storage, Parallel Jobs, etc. This justifies the increase of the percentage of the development effort in the last two years of the project.

	Involved Partners	PY1	PY2	PY3 (excluding PQ12)
GOCDB	STFC	5	5	5
Ops Portal	CNRS	17	20	16
GGUS	KIT-G	10	13	9
SAM (including messaging support)	CERN SRCE GRNET	13	11	6
Accounting Repository	STFC	1,6	7	5
Accounting Portal	FCTSG/CSIC	6	3	4
Metrics Portal	FCTSG/CSIC	0,6	2	1

Table 1 – Reported effort tool by tool for the first three years of the project, excluding PQ12.

Tool	PY1		PY2		PY3		PY4 Forecast	
	Dev	Maint	Dev	Maint	Dev	Maint	Dev	Maint
GOCDB	50%	50%	50%	50%	50%	50%	50%	50%
Ops Portal	80%	20%	80%	20%	55%	45%	50%	50%
GGUS	55%	45%	55%	45%	50%	50%	50%	50%
SAM	45%	55%	35%	65%	30%	70%	30%	70%
Accounting Repository	75%	25%	75%	25%	75%	25%	50%	50%
Accounting Portal	20%	80%	30%	70%	80%	20%	80%	20%
Metrics Portal	90%	10%	90%	10%	80%	20%	80%	20%

Table 2 - Maintenance versus Development Effort, tool by tool, for the first three years of the project and the forecast for the fourth.

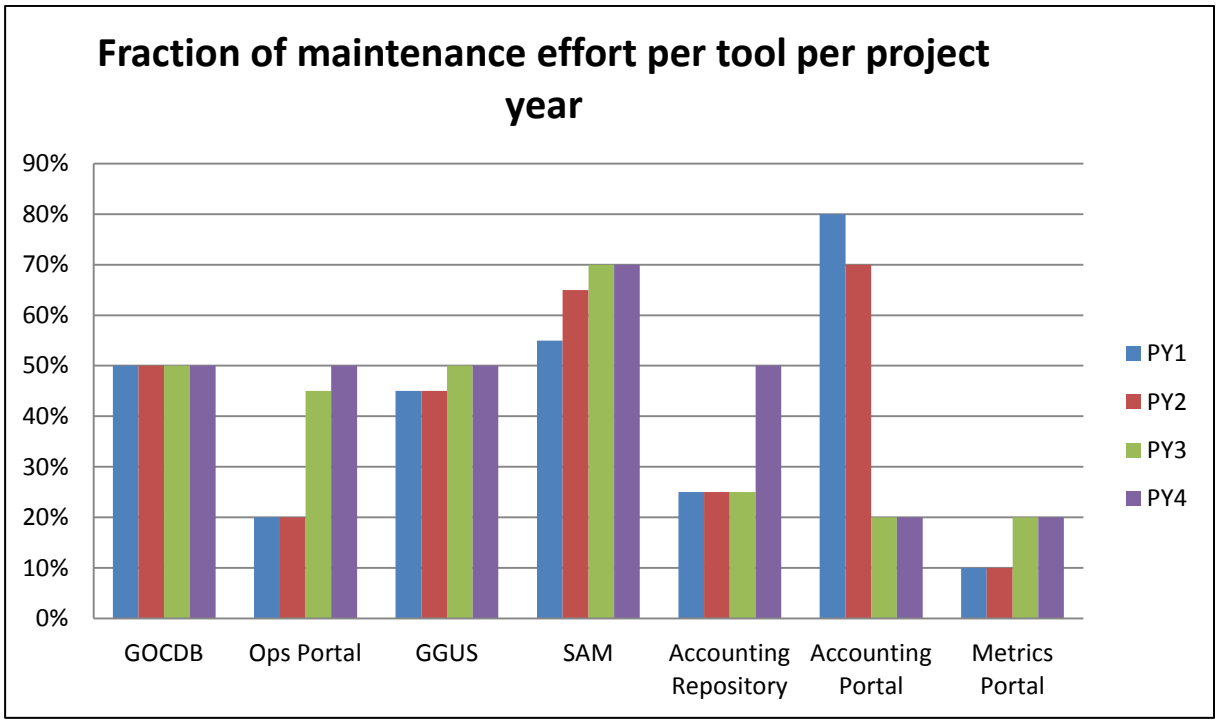

Figure 6 - Evolution of the fraction of effort devoted to maintenance tool by tool for the four year of the project. The point for PY4 is a forecast.

4.3 Major Issues

The following major issues have been encountered and addressed during PY3:

- **Effort for new developments in SAM**: The effort available in the SAM team for new development is very limited taking into account the complexity reached by the system. Some important requirements (viz. support for virtual sites, monitoring non-EGI sites and services) have been rejected by the product team or are still in waiting. All the effort will be mainly spent on maintenance and 3rd level support.

  Mitigation: Where possible new developments will be moved to other operational tools. For example, a mini-project to monitor non-EGI sites and services has been funded by EGI-InSPIRE in the context of the call for mini-projects of January 2013. This activity will implement the basis of a new A/R reporting service using open source project components that will be more flexible and extensible and allow the inclusion of more middleware services into the calculation of A/R metrics and will also add VO-wise metric results (in addition to service-wise, site-wise and NGI-wise provisioning of results). The project is led by GRNET and includes IN2P3 (CNRS) and SRCE.

- **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 Grid users many of whom may have only a 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 (viz. from EMI and IGE), also increased the number of GGUS users outside EGI. This means that the EGI's governance of operational tools (OMB and OTAG) may lack

representation of important GGUS user communities and thus would not be the most suitable place to discuss GGUS requirements prioritization.

Mitigation: A new GGUS advisory board has been established during the EGI Technical Forum in Prague. It is composed of representatives from user communities, NGIs, EGI, technology providers. Meeting frequency is once per month directly after the monthly GGUS releases. The new requirements will be discussed there before reaching the OTAG.

- **Single access point to releases information**: as extensively discussed in D7.1 at the beginning of the project it was decided to organise the JRA1 activity as a set of independent product teams with a central coordination. This means that each PT has its own bug tracking tool, build system, code repository and documentation pages. The requirements tracking tool was unified according to the project guidelines (with the exception for GGUS discussed in the previous point) 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 keeping the wiki pages up to date, it is still difficult to get information about the schedule and content of future releases. This information exists for each tool of course, but a single aggregation point that shows in just one page what is going on within JRA1 is still missing.

  Mitigation: A dedicated section with all the links to release notes and schedule for all the tools has been created on the EGI wiki [R 60]. However, the information is still not updated regularly. A solution to get and publish this information in a (semi)automatic way will be investigated. The usage of the "ops-tools-roadmap" RT queue is now abandoned.

- **Integration of the EMI probes in SAM**: the development of the probes is out of scope of SAM and JRA1 activities. Several probes have been developed in the context of the EMI project and the JRA1 team should integrate them into the SAM framework. Some problems have been encountered concerning dependencies, environments and configurations in the definition of the packages that EMI should provide to JRA1. Moreover, the roles of EMI and the JRA1 teams in the integration process were not well defined. This caused some delay in the probe integration.

  Mitigation: A meeting between EMI and JRA1 teams was organized to discuss the aforementioned problems[5]. The main outcomes are that JRA1 provided EMI with a set of requirements that should be fulfilled by the EMI packages and the definition of the roles for both teams. EMI are responsible for testing the package contents, while JRA1 tests the integration in SAM and is responsible for that integration. We decided to create an integration test-bed [R 61] to simplify the communication between the teams during the integration process. The status of the integration process and the main issues identified by the JRA1 team, that have to be fixed by EMI, are reported there. Moreover, the EGI SAM probes working group, composed of experts from NGIs, EMI and EGI, was created to revise the EMI probes before they are integrated in the SAM framework and to evaluate probe- and monitoring-related improvements.

[5] https://indico.egi.eu/indico/conferenceDisplay.py?confId=1088

5 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 MS710 "Roadmap for the Maintenance and Development of the Deployed Operational Tools" [R 62] which details the roadmap until the end of the project. However a short list of major expected developments in PY4 for each tool is provided in the following tables.

Operations Portal

	PQ13	PQ14	PQ15	PQ16
Availabilities and reliabilities report system	X			
Refactoring of the different dashboards	X			
Continuous integration			X	
Package				X

Table 3 - Operations Portal roadmap for PY4.

GGUS

	PQ13	PQ14	PQ15	PQ16	TBD
Integration of the last remaining NGI (Russia)	X				
GGUS Report Generator (final version depends on external requirements)	X				
Adapt interface to GOC DB/Doctrine	X				
Implementation of alarm processes for EGI tools	X				
High availability for GGUS components (switching between stacks)		X			
High availability for all GGUS interfaces			X		
Disaster recovery plan			X		
Additional authentication through shibboleth			X		
New interfaces to PRACE, MAPPER, DANTE					X
Specific work flows for CSIRT/Security					X

Tabella 4 - GGUS roadmap for PY4.

GOCDB

	PQ13	PQ14	PQ15	PQ16
GOCDB v5	X			

Extending Scoping		X	X	
Extensibility Mechanism			X	X
Glue2 XML Rendering and add Glue2 Downtime			X	X

Table 5 - GOCDB roadmap for PY4.

SAM

	PQ13	PQ14	PQ15	PQ16
Integration of some EMI probes		X		
Messaging: Implementation of SAM probes failover capabilities			X	

Table 6 - SAM roadmap for PY4.

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

	PQ13	PQ14	PQ15	PQ16
EMI-APEL 3 Client released[6]	X			
Regional APEL Server released[7]	X			
Cloud Accounting Summaries to Accounting Portal[8]	X			
Application Accounting usage record defined[9]	X			
Storage Accounting Summaries to Accounting Portal		X		
Publishing summaries from Accounting Repository to other sites (OSG/DGAS)			X	

Table 7 - Accounting Repository roadmap for PY4.

Accounting Portal

	PQ13	PQ14	PQ15	PQ16
EGI User usage accounting[10]	X			
Provisioning of Cloud Accounting (Preliminary support) [11]	X			
Regional portal codebase improvements[12]	X			

[6] First release foreseen in PQ12.
[7] First release foreseen in PQ12.
[8] First release foreseen in PQ12.
[9] First release foreseen in PQ12.
[10] First release foreseen in PQ12.
[11] First release foreseen in PQ12.
[12] First release foreseen in PQ12.

XML endpoints generalization and improvement[13]	X			
EGI Usage VT Report Improvements[14]	X			
EGI Usage VT Publishing Improvements[15]	X			
Scientific Disciplines VT Interface Support	X			
Scientific Disciplines VT final Implementation	X			
SSM implementation for normal Accounting	X			
Provisioning of Cloud Accounting (Improvements)		X		
Contributed CPUs by site		X		
Preliminary support for parallel (MPI) jobs		X		
Provisioning of Storage accounting (DB implementation)		X		
Provisioning on Storage accounting (View Implementation)		X		
Provisioning of MPI accounting (DB implementation)		X		
Provisioning of MPI accounting (View Implementation)		X		
Provisioning of Application accounting (DB implementation)		X		
Provisioning of Application accounting (View implementation)		X		
Revised version of all the new features				X

Table 8 - Accounting Portal roadmap for PY4.

Metrics Portal

	PQ13	PQ14	PQ15	PQ16
Manual metrics expansion and refinement		X		
New customized reports with Excel support		X		
Views enhancement and optimization		X		
GGUS metrics improvement and new A/R metrics		X		
Revised version of all the new features				X

Table 9 - Metrics Portal roadmap for PY4.

[13] First release foreseen in PQ12.
[14] First release foreseen in PQ12.
[15] First release foreseen in PQ12.

6 PARTICIPATION IN CONFERENCES AND PROJECT TASK FORCES

JRA1 and all its product teams actively attended the annual main project events, the EGI Technical Forum and the EGI Community Forum. During the Technical Forum in Prague two workshops were organized: one on accounting and the other on future advancements of tools and regionalisation. In the accounting workshop several strands of work in EGI, EMI, and other projects to define and collect usage records for storage, virtual machine/cloud use, MPI jobs, and applications were described. The workshop on future advancements of tools and regionalisation analysed and discussed the foreseen roadmap for the tools, in particular the evolution of those that are (or will be) regionalised to be deployed at the NGI level. In the EGI Community Forum which will be held in Manchester in April 2013, two more workshops will be organized: the first again on accounting while the second will discuss the tools evolution and the requirements to be fulfilled to allow usage in other infrastructures like EUDAT and PRACE.

JRA1 organized the Long Term Sustainability of Operational and Security Tools workshop in Karlsruhe, Germany [R 63] in September to discuss how to maintain the operational tools after the project.

The GOCDB team attended OGF meetings to participate in the GLUE2 working group.

JRA1 has been represented in all the project taskforces, groups and bodies that required its presence, including the new GGUS Advisory Board. Finally, one OTAG meeting was held during PY3.

7 CONCLUSIONS

This document reported the third year activity of the EGI-InSPIRE 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 the 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. TJRA1.5 for the development needed for the extension of the Operations Portal and its harmonization with other portal frameworks

TJRA1.5 completes its activities at the end of this year, while TJRA1.3 ended in PY2.

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:

- **Operations Portal:**
  - an important dashboards refactoring and a new dashboard dedicated to VO Operations;
  - new modules to compute and expose the availabilities and reliabilities of the TOP-BDII of a NGI and sites;
  - a new module for computing VO specific availabilities/reliabilities reports;
- **GGUS:**
  - the Report Generator has been redesigned from scratch;
  - High Availability (HA) solution;
  - a new GGUS advisory board has been established. It is composed of representatives from user communities, NGIs, EGI, technology providers.
- **GOCDB:**
  - a separate GOCDB instance has been deployed by the EUDAT project;
  - the design and development of GOCDB v5 started. It is based on a new data layer able to use different RDBMS platforms (e.g. MySQL, Postgres, Oracle);
- **Accounting system:**
  - the new SSM (Secure Stomp Messenger) based APEL Accounting Repository has been brought in production in June 2012;
  - the regional accounting repository and the regional accounting portal will be deployed in production, respectively, in April and May 2013;
- **SAM:**
  - POEM component has been included and fully integrated in SAM Update-17;
  - MyEGI has been reviewed and improved as part of SAM Update-19;
  - SAM instance for monitoring operational tools (OPS-MONITOR) has been deployed as part of SAM Update-20;
- **Metrics Portal:**
  - new metrics types added.

For what concerns the integration of new (not gLite) middleware types into the operational tools, this is now complete from the perspective of SAM and GOCDB for:

- UNICORE
- ARC
- GLOBUS

- Desktop Grids
- QosCosGrid

In addition, operational tools and other EGI services, like the Applications Database, the Training Marketplace and the Customer Relation Manager are now monitored by SAM.

In PY3 **regionalisation activities** have been carried on under the umbrella of TJRA1.2. The product teams performed a series of analyses to understand the benefits of having a regionalised version, taking into account the real NGI needs, the possibility to create regional views on the central instance and the effort needed to develop and maintain the regionalised package. After several discussions with NGIs, during the OTAG meetings or by surveys, we decided to drop the Operations Portal and GOCDB regional versions. NGIs agreed to use regional views on the central instance. Regarding accounting, the development continued during PY3 and the regionalised releases of the Accounting Repository and Portal will be deployed in production, respectively, in April and May 2013.

After a first year focused on the investigation, **TJRA1.4** mainly spent PY3 to develop the solutions defined at the end of PY2. Cloud, CPU and Parallel Jobs resource types now can be accounted in the new EGI Accounting Repository based on SSM and steps have been taken to have accounting data for applications, storage and virtualisations. The relevance of the inclusion of the new resource types in the accounting system for the EGI users and, in particular, the steadily increasing importance of the Clouds world, pushed JRA1 to mainly devote TJRA1.4 effort to these activities.

TJRA1.5 activities were reported in section 2.1 dedicated to the Operations Portal. This task completes its activities in PY3 after the development of several major features during the first three years of the project. During PY4, the Operations Portal team will be devoted mainly to maintenance activities, in the context of TJRA1.2.

The main issues that arose during the year are reported in section 4.

Section 4 also provides a spent **effort assessment** for all the tools giving an estimate of the splitting between the effort needed for base maintenance and for new feature 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 has grown for almost all tools during the project lifetime. This reflects the decrease of the number of new features to be developed once several user requirements have already been fulfilled and the increase of the effort needed to refine the already delivered capabilities (e.g. bug fixing). The maintenance effort will start to decrease when the new functionalities become more stable and the features set quite frozen. Exceptions to this trend are the GOCDB and the Accounting Portal due to the development of GOCDB v5 and the integration of the new resource types in the Accounting Portal.

A single OTAG meeting was organized during PY3 to assess the decrease of the number of new requirements received. This could be considered as proof that **the tool consolidation phase has begun**. Discussions on minor requirements continued in RT and OMB mailing lists and during JRA1 meetings.

The PY4 roadmap for the tools is not described in detail in this document; however, it was briefly outlined in section 5 as it was the focus of a recent project milestone (MS710).

8 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	SA1 Main Wiki Page: https://wiki.egi.eu/wiki/EGI-inSPIRE_SA1
R 4	Top-BDII Availability and Reliability: https://wiki.egi.eu/wiki/Top-BDII_list_for_NGI#top-BDII_Availability_and_Reliability
R 5	VO monthly A/R statistics - requirement: https://rt.egi.eu/rt/Ticket/Display.html?id=4348
R 6	VO monthly A/R statistics – Design document: https://documents.egi.eu/secure/ShowDocument?docid=1493
R 7	Operations Portal Release Notes Browser: http://operations-portal.egi.eu/aboutportal/releaseNotesBrowser
R 8	MS410 EGI Helpdesk and the NGI support units: https://documents.egi.eu/document/522
R 9	GGUS advisory board: https://indico.egi.eu/indico/sessionDisplay.py?sessionId=58&confId=1019#20120921
R 10	PRACE: http://www.prace-ri.eu/
R 11	DANTE (Delivery of Advanced Network Technology to Europe): http://www.dante.net/
R 12	CSIRT: https://wiki.egi.eu/wiki/EGI_CSIRT:Main_Page
R 13	xGUS helpdesk: http://xgus.scc.kit.edu/
R 14	GOCDB change log https://www.sysadmin.hep.ac.uk/svn/grid-monitoring/tags/gocdb/GOCDB-4.4/changeLog.txt
R 15	GOCDB failover instance https://goc.itwm.fraunhofer.de/portal/
R 16	EUDAT project: http://www.eudat.eu/
R 17	EGI-InSPIRE DoW: https://documents.egi.eu/document/1061
R 18	Call for Projects to Accelerate EGI's Strategic Objectives: https://indico.egi.eu/indico/getFile.py/access?contribId=0&resId=0&materialId=0&confId=1100
R 19	MS 706 Operational Tools Accounting Work Plan: https://documents.egi.eu/document/531

R 20	Secure Stomp Messenger: https://wiki.egi.eu/wiki/APEL/SSMOverview
R 21	IGE Project Homepage: http://www.ige-project.eu/
R 22	QoSCoSGrid Middleaware: http://www.qoscosgrid.org
R 23	European Desktop Grid Initiative: http://edgi-project.eu/
R 24	Unicore: http://www.unicore.eu/
R 25	Central SAM documentation: https://tomtools.cern.ch/confluence/display/SAMDOC/Home
R 26	SAM Update-17 Release notes: https://tomtools.cern.ch/confluence/display/SAMDOC/Update-17.1
R 27	SAM Update-19 Release notes: https://tomtools.cern.ch/confluence/display/SAMDOC/Update-19
R 28	SAM Update-20 Release notes: https://tomtools.cern.ch/confluence/display/SAMDOC/Update-20
R 29	Integration of EMI probes in SAM: https://twiki.cern.ch/twiki/bin/view/EMI/NagiosServerEMITestbed0022012
R 30	EGI SAM probes Working Group: https://wiki.egi.eu/wiki/Nagios_probes_Working_Group
R 31	Local SAM regional instances: https://wiki.egi.eu/wiki/SAM_Instances
R 32	EGI software provisioning process: https://wiki.egi.eu/wiki/EGI_Software_Provisioning
R 33	ActiveMQ message system homepage: http://activemq.apache.org/
R 34	GOCDB Service Types Query: https://goc.egi.eu/gocdbpi/public/?method=get_service_types
R 35	iRods: https://www.irods.org/
R 36	ARC: http://www.nordugrid.org/arc/
R 37	BES: http://www.ogf.org/UnderstandingGrids/documents/OGSA-BES-OnePager%20jfe.pdf
R 38	Torque: http://www.adaptivecomputing.com/products/open-source/torque/
R 39	Squid: www.squid-cache.org
R 40	XRootD: http://xrootd.slac.stanford.edu/

R 41	GOCDB documentation – Adding new services types: https://wiki.egi.eu/wiki/GOCDB/Input_System_User_Documentation#Adding_new_services_types
R 42	DMSU Main Wiki Page: https://wiki.egi.eu/wiki/TSA2.5_Deployed_Middleware_Support_Unit
R 43	EMI Project Homepage: http://www.eu-emi.eu/
R 44	SAGA: www.ogf.org/documents/GFD.90.pdf
R 45	FedCloud Main Wiki Page: https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce
R 46	GGUS documentation – Adding a new support unit: https://wiki.egi.eu/wiki/FAQ_GGUS-New-Support-Unit
R 47	FedCloud Scenario5 Wiki Page: https://wiki.egi.eu/wiki/Fedcloud-tf:WorkGroups:Scenario5
R 48	FedCloud Scenario4 Wiki Page: https://wiki.egi.eu/wiki/Fedcloud-tf:WorkGroups:Scenario4
R 49	TJRA1.4 wiki page: https://wiki.egi.eu/wiki/Accounting
R 50	Inter Project Accounting TaskForce https://wiki.egi.eu/wiki/TCB:Accounting_Task_Force
R 51	Evolving EGI Workshop: http://indico.egi.eu/indico/conferenceDisplay.py?ovw=True&confId=1252
R 52	EGI Pay-for-Use Pilot Group: https://wiki.egi.eu/wiki/EGI_Pay-for-Use_Pilot
R 53	OpenNebula: http://opennebula.org/
R 54	OpenStack: http://www.openstack.org/
R 55	WNoDes: http://web2.infn.it/wnodes/index.php/wnodes
R 56	Okeanos: https://okeanos.grnet.gr/home/
R 57	OGF Usage Record Working Group: http://www.ogf.org/gf/group_info/view.php?group=ur-wg
R 58	StaR – Storage Accounting Record: http://bit.ly/PoYu3x
R 59	D7.1- Annual Report on Operational Tool maintenance and development activity https://documents.egi.eu/document/372

R 60	Operational Tools testbed: https://wiki.egi.eu/wiki/Operational_tools_testbed
R 61	Integration of EMI probes in SAM – Integration testbed: https://twiki.cern.ch/twiki/bin/view/EMI/NagiosServerEMITestbed0022012
R 62	MS710 - Roadmap For The Maintenance And Development Of The Deployed Operational Tools: https://documents.egi.eu/secure/ShowDocument?docid=1501
R 63	Long Term Sustainability of Operational and Security Tools workshop in Karlsruhe, Germany, September 2013: https://indico.egi.eu/indico/conferenceDisplay.py?confId=1132

9 APPENDIX A: JRA1 EFFORT TABLES

		Project Period 1		
Task	Partner	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%

		Project Period 2		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.2	10B-KIT-G	13,4	11,8	114%
TJRA1.2	12B-FCTSG-CSIC	4,4	3,0	146%
TJRA1.2	14A-CNRS	3,0	3,0	100%
TJRA1.2	16A-GRNET	1,9	3,0	63%
TJRA1.2	17-SRCE	3,5	3,0	116%
TJRA1.2	34A-STFC	5,3	6,0	88%
TJRA1.2	35-CERN	5,7	3,0	189%
Total:		37,1	32,8	113%

		Project Period 3 (QR9-QR11)		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.2	10B-KIT-G	9,0	8,8	102%
TJRA1.2	12B-FCTSG-CSIC	1,9	2,3	84%
TJRA1.2	14A-CNRS	2,1	2,3	93%

TJRA1.2	16A-GRNET	1,6	2,3	71%
TJRA1.2	17-SRCE	2,5	2,3	111%
TJRA1.2	34A-STFC	4,4	4,5	98%
TJRA1.2	35-CERN	2,1	2,3	93%
	Total:	23,6	24,6	96%

		PY1+PY2+PY3		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.2	10B-KIT-G	32,0	32,3	99%
TJRA1.2	12B-FCTSG-CSIC	11,9	8,3	144%
TJRA1.2	14A-CNRS	8,2	8,3	99%
TJRA1.2	16A-GRNET	4,4	8,3	54%
TJRA1.2	17-SRCE	9,3	8,3	112%
TJRA1.2	34A-STFC	14,9	16,5	91%
TJRA1.2	35-CERN	8,3	8,3	101%
	Total:	88,9	90,1	99%

Table 10 - Effort for TJRA1.2 in PY1, PY2 and PY3 excluding PQ12.

		PY1+PY2		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.3	12B-FCTSG-CSIC	2,9	3,0	96%
TJRA1.3	14A-CNRS	3,5	3,0	116%
TJRA1.3	17-SRCE	3,6	3,0	119%
TJRA1.3	34A-STFC	2,9	3,0	96%
TJRA1.3	35-CERN	4,3	6,0	72%
	Total:	17,1	18,0	95%

Table 11 - Effort for TJRA1.3 in PY1 and PY2.

		Project Period 2		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.4	10H-LUH	6,6	6,0	111%
TJRA1.4	12B-FCTSG-CSIC	0,7	6,0	12%
TJRA1.4	21A-INFN	1,2	8,7	14%
TJRA1.4	34A-STFC	6,9	9,0	77%
	Total:	15,5	29,7	52%

		Project Period 3 (QR9-QR11)		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.4	10H-LUH	3,3	4,5	73%
TJRA1.4	12B-FCTSG-CSIC	2,7	4,5	60%
TJRA1.4	21A-INFN	1,9	6,5	29%
TJRA1.4	34A-STFC	4,5	6,8	67%
	Total:	12,4	22,3	56%

		PY2+PY3		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.4	10H-LUH	9,9	10,5	95%
TJRA1.4	12B-FCTSG-CSIC	3,4	10,5	33%
TJRA1.4	21A-INFN	3,1	15,2	21%
TJRA1.4	34A-STFC	11,4	15,8	72%
	Total:	27,9	51,9	54%

Table 12 - Effort for TJRA1.4 in PY2 and PY3 excluding PQ12.

		Project Period 1		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	13,488	17,667	76%

		Project Period 2		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	17,5	17,7	99%

		Project Period 3 (QR9-QR11)		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	14,3	13,3	108%

		PY1+PY2+PY3		
Task	Partner	Worked PM Funded	Committed PM	Achieved PM %
TJRA1.5	14A-CNRS	45,3	48,6	93%

Table 13 - Effort for TJRA1.5 in PY1, PY2 and PY3 excluding PQ12.