



EGI-InSPIRE

ANNUAL REPORT ON OPERATIONAL TOOLS MAINTENANCE AND DEVELOPMENT ACTIVITY

EU DELIVERABLE: D7.1

Document identifier:	EGI-372-V5.doc
Date:	23/05/2011
Activity:	JRA1
Lead Partner:	INFN
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/372



Abstract

The Joint Research Activity1 (JRA1 or Work Package7) is responsible for developing and maintaining operational tools for the EGI-InSPIRE project. The list of tools includes: the topology and static information database (GOCDDB), the Operation Portal providing a single access point to operational information gathered from various sources, the EGI Helpdesk based on the GGUS central system interfaced with regional helpdesks, the SAM monitoring framework that detects failures at sites and services, the Accounting Repository that stores accounting data, its web interface (the Accounting Portal) and the Metrics Portal which allows the insertion, storage and display of the project's performance metrics. Additional goals of the activity are supporting the configuration of the production message broker network and the extension of the accounting system to encompass the new resource and middleware types that will be included in the production infrastructure. The evolution of the accounting towards an economic and billing system will be taken into account by the activity as a fundamental step to reach the EGI's self-sustainability. The work performed to maintain and evolve the tools and to accomplish activity tasks during the first year of the project is described in this report. The activity's major achievements and issues faced during the last year are identified and discussed.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. 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. 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	D.Cesini	INFN/JRA1	25/03/2011
Reviewed by	Moderator: Tiziana Ferrari Reviewers: Jeremy Coles Rebecca Breu	EGI.eu STFC JUELICH	18/04/2011
Approved by	AMB & PMB		23/05/2011

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	25/03/2011	First draft (V1)	D.Cesini/INFN
2	27/03/2011	V2 according to Newhouse review	S.Newhouse / D.Cesini
3	27/04/2011	V3 ready for second round of reviews	D.Cesini
4	11/05/2011	V4 ready after second round	D.Cesini

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

The JRA1 activity is responsible for developing and maintaining operational tools for the EGI-InSPIRE project. The list of tools includes: the topology and static information database (GOCDB), the Operation Portal providing a single access point to operational information gathered from various sources, the EGI Helpdesk based on the Global Grid User Support (GGUS) central system interfaced with regional helpdesks, the Service Availability Monitor (SAM) monitoring framework that detects failures at sites and services, the Accounting Repository that stores accounting data, its web interface (the Accounting Portal) and the Metrics Portal which allows the project's performance metrics to be inserted, stored and displayed.

Additional goals of the activity are supporting the configuration of the production message broker network and the extension of the accounting system to encompass the new resource and middleware types that will be included in the production infrastructure. The evolution of the accounting system towards an economic and billing model will be taken into account by the activity as a fundamental step to allowing EGI to reach self-sustainability. This document will present how, at the beginning of the project, the activity was set up following a model of centrally coordinated independent, geographically distributed development teams and how a common requirements workflow was put in place.

The main achievements over the last year in the development and maintenance activity are presented with details for all of the tools, these include: i) the release to production of GOCDB4, which constituted an important step towards the regionalisation of the tool; ii) the refactoring of important SAM monitoring components with the release of the new myEGI visualization tool and the inclusion of ARC middleware probes; iii) a complete refactoring and regionalisation of the Operations Portal with the inclusion of the security dashboard; iv) the renaming and restructuring of GGUS support units and tickets workflow to accommodate the new EGI structure and the release of the xGUS system; v) the migration of the accounting system to the messaging infrastructure for the data transport layer; vi) the development of a completely new metrics portal and vii) the migration to the egi.eu domain for all the operational tools and the rebranding of the web pages to show EGI logos. Achievements and issues on the regionalisation for all the tools are also presented in the document. Discussions on regionalisation use cases and models are still controversial and a definite development plan is not ready for all the tools.

Eight main sections compose the document: Section 2 briefly presents the developed tools and describes the actions taken at the beginning of the project to setup the activity and to ease the transition from the previous EGEE-III project (where a dedicated workpackage for operational tools development was not present and the needed effort was embedded in the operation workpackage); Section 3 is a report on the maintenance work and the new feature implementation for all the tools; Section 4 provides details on the transition towards a regionalised model for some of the tools – this work started during the EGEEIII project but only the SAM tool was already fully regionalised at the beginning of EGI-InSPIRE; Section 5 is dedicated to the Operations Portal on its harmonization and integration with other portals and tools. The main issues from the first year, a conclusion and summary of the organised/attended meetings are discussed in the remaining sections.



TABLE OF CONTENTS

1	INTRODUCTION	7
2	JRA1 OVERVIEW	8
2.1	Operational tools	8
2.2	Product Teams and the development infrastructure	9
2.3	The requirements workflow	10
3	DEVELOPMENT AND MAINTENANCE	12
3.1	Operations Portal	12
3.2	EGI Helpdesk (GGUS)	13
3.3	Grid Configuration repository (GOCDB)	14
3.4	Accounting Repository	15
3.5	Accounting Portal	16
3.6	Service Availability Monitoring	16
3.7	Metrics Portal	18
3.8	Broker network configuration and infrastructure	18
3.9	Ops Tools migration to egi.eu domain	18
3.10	Integration of new middleware types	19
3.11	Monitoring tools availability	19
3.12	Dependencies and interaction with externally developed tools	19
4	REGIONAL MODEL FOR OPERATIONAL TOOLS	21
4.1	Background	21
4.2	Operations Portal and Dashboard	21
4.3	EGI Helpdesk	22
4.4	Grid Configuration repository (GOCDB)	23
4.5	Accounting Repository	23
4.6	Accounting Portal	25
4.7	Service Availability Monitoring	25
4.8	Metrics Portal	25
5	ACCOUNTING FOR DIFFERENT RESOURCE TYPES	26
6	INTEGRATED OPERATIONS PORTAL	27
6.1	Porting to the Symfony open-source web-development framework	27
6.2	Towards a service oriented model	27
6.3	Interoperability with GOCDB	27
6.4	Changes to the initial workplan	29
6.5	Security Dashboard	30
6.6	COD View	30
7	FIRST YEAR ISSUES AND MITIGATION	31
8	ATTENDED/ORGANISED CONFERENCES AND MEETINGS	33
9	CONCLUSION	34
10	REFERENCES	35



1 INTRODUCTION

The Joint Research Activity1 (JRA1) is responsible for the development and maintenance of the operational tools used by the European Grid Infrastructure. Most of the tools and development teams were inherited from past EGEE projects but of course development is still needed to maintain the software and to follow the evolution of the infrastructure that is expected to encompass additional resource types, such as desktop grids, cloud and high capacity computing resources. Moreover the much more distributed model for the project operations must be reflected into the development and deployment of the tools. The work started during the EGEE-III project to allow for a fully distributed or regionalized model of the operational tools will be completed. The evolution of the accounting system towards an economic and business model is also an objective of the activity.

Five tasks compose the activity:

- TJRA1.1 – Activity management: The coordination of the development work, including the definition of the release roadmaps, the representation of the activity within the projects bodies and the reporting on the status of the activity. This is a 4 year task.
- TJRA1.2 – Maintenance and development of the deployed operational tools: The underlying bug fixing and development work for the operational tools. This is a 4 year task.
- TJRA1.3 – National Deployment Models: Development needed to support the deployment at a regional (or NCI) level. This task is expected to be completed by the end of the first year of the project.
- 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. This is a 3 year task starting from the second year of the project.
- TJRA1.5 – Integrated Operations Portal: The development needed for a restructuring of the Operations Portal and its harmonization with other portal frameworks such as the Grid Configuration Database (GOCDB). This task spans over the first three years of the project.

The main preliminary steps of the Joint Research Activity consisted in the setting up of the development infrastructure and of the release procedures and in the definition of the requirements gathering workflow to establish proper communication channels. After a brief description of all the JRA1 tools, the subsections that follow describe how this was achieved. Further details on the establishment of the production infrastructure are available in the project milestone MS702 [R1].



2 JRA1 OVERVIEW

2.1 Operational tools

The EGI-InSPIRE JRA1 activity is responsible for the development and maintenance of the following operational tools:

- **Operations Portal:** (the former CIC Portal) a single access point to operational information. This is used by all the project actors¹. It is composed of various modules the main ones being:
 - o The Broadcast tool that allows users to send bulk messages to various communities
 - o The Operational Dashboard that collects information from many sources about site/service status and failures and allows users to semi-automatically open tickets in the EGI Helpdesk system
 - o A VO Identity Cards repository that stores and provides an interface for inputting Virtual Organisation (VO) static information for wider reference.
- **EGI Helpdesk:** the main support access point for the project. Primarily used by Grid end-users and by support teams, but also by site managers who need support to solve middleware issues. It is based on the central Global Grid User Support (GGUS) system interfaced where possible with regional helpdesks.
- **Grid Configuration Database (GOcdb):** contains general and semi-static information about the sites participating in the production Grid (covering data such as site services installed, site manager contact details, security contacts etc.). It is accessed by all project actors, by other tools and by third party middleware in order to obtain an operational snapshot or a specific bit of information.
- **Accounting Repository:** stores data about VO usage of site resources within the production infrastructure. Primarily accessed by other tools (e.g. the Accounting Portal) in order to create usage reports. It is based on the gLite-APEL system [R21]. APEL sensors that provide data to be stored in the Accounting Repository are not developed by JRA1, but within the EMI project.
- **Accounting Portal:** graphical frontend for the Accounting Repository. It is accessible to anyone with a recognised grid certificate and displays overviews with details determined by the role associated with the certificate (for example user, site administrator, regional manager and VO manager).
- **Service Availability Monitoring (SAM):** a monitoring framework for sites and services. Used by site managers to monitor their sites, but also at a project level as a data source to create availability/reliability statistics for the resource centers participating in the production infrastructure. It is one of the main sources of data for the Operations Dashboard described previously. It is composed of various components, the most important being:
 - o The test submitting framework: based on the NAGIOS system [R49] set up and customized by the NAGIOS Configurator (NCG)
 - o The DataBase components: The Aggregated Topology Provider (ATP) [R24], the Metric Description DataBase (Mddb) and the Metrics Result DataBase (MRDB)
 - o A message bus to publish the monitoring results

¹ Throughout the document the term project actors refer to end-users, site managers, NGI managers, support teams and VO managers.



- A visualization tool called MyEGI.
- **Metrics Portal:** collects a set of metrics from different resources to help in measuring project performance and keep track of the project evolution by displaying historical values of the metrics in a single place. It also provides web interfaces to inject the metrics into the database.

The JRA1 activity is also responsible for providing support to the configuration of the message broker network of the production infrastructure based on the ActiveMQ system [R50].

An activity task (TJRA1.4), starting in the second year of the project, will also add functionalities to the accounting system in order to account for more resource types, details will be given in Section 5.

2.2 Product Teams and the development infrastructure

As anticipated most of the tools and development teams were inherited from the previous EGEE projects where the development work was carried out within the operation work package (SA1). In EGI-InSPIRE the operational tools development was moved to a standalone activity (JRA1). In the following we will indicate development teams within JRA1 as Operational Tools Product Teams (OTPTs).

The inherited OTPTs were geographically distributed across Europe and used different development infrastructures, in terms of bug/task tracking tools, repositories, building tools, documentation pages etc, so during the early stages of the activity one of the main tasks was the decision on how to evolve the OTPTs and if it was needed to merge them into a single development infrastructure. It was decided not to change the local, independent infrastructures, but to agree on a common release procedures and requirements workflows.

This decision had the advantage that all the OTPTs were immediately operative; they could release new versions of their tools from the very beginning of the project and did not lose time in changing their development facilities and habits. The drawback was that for actors external to the activity it was sometimes difficult to obtain a global view of the ongoing work within JRA1. To address this situation a section dedicated to JRA1 in the EGI wiki [R2] was used to provide links to the relevant distributed information, in particular regarding documentation resources that are maintained by OTPTs on their local infrastructure. JRA1 tries to maintain these wiki pages as much as possible. The EGI Request Tracker (RT) system [R5] is also used to provide an overview of the future JRA1 releases.

A detailed description of all the OTPTs, their composition, location and development facilities, together with an overview of the common release procedures can be found in the project milestone MS702 [R1], while the common requirements workflow is presented in the next section.

During the early stages of the project it was agreed with the SA2 and SA1 activities treat JRA1 products that need to be deployed at the regional level as any other software or middleware product that is installed in the production infrastructure and that is generally provided by third parties or by other projects (e.g. EMI). This implies that JRA1 tools that need deployment must successfully undergo the quality criteria verification, repository mirroring and staged rollout steps foreseen by the SA1 and SA2 activities and described in MS402 [R3] and MS503 [R4]. During the first year of the projects only the SAM tool, being the only one deployed in all NGIs, underwent the official project release process. Future releases of the Regional Operations Portal will probably follow the process, but it was decided to avoid it during the first year given the small number of NGIs that deployed it. Details on the regionalisation status for the tools will be given in Section 4.

Interactions between the OTPTs are encouraged by the activity management through periodic phone conferences and face-to-face meetings. Activity progress is tracked using EGI.eu facilities such as the



RT system [R5], and is reported in periodic documents. A dedicated mailing list (under the egi.eu domain) is also available for activity communications.

2.3 The requirements workflow

In this section and throughout the whole document the word “requirements” refers to new feature requests or to requests about changes in behaviour of some components. It does not refer to bug fixing requests. As described in this section JRA1 identified a common (to all development teams) requirements workflow, while bugs are handled and tracked independently by the OTPTs.

The common requirement workflow definition for all the tools is now complete, even if it took more time than expected, in fact JRA1 started with its own workflow using a dedicated RT queue to track development requests, but it was then agreed with other activities that need to handle requirements (mostly NA3 and SA1) to make the processes uniform and to use the same tracking tool, i.e. the same RT queue properly configured with custom fields to allow for a convenient distribution of requirements among the activities. In the following, this RT queue will be referred to as the “Requirement queue”.

Users of the JRA1 tools can create requirements anytime by filing a ticket in the Requirement queue and addressing it to the “Operational Tools” category. Requirements that arrive through other ways (e.g. emails, phone conferences, documents etc.) are also translated into RT tickets. Periodically (every two months at most) each JRA1 PT reviews the requirements list accepting those that can be addressed immediately because they do not require much effort or do not break any interaction with other tools and labels all the others items as “to be discussed”. Discussions and prioritization of those requirements take place during internal JRA1 meetings at first and then within the Operational Tools Advisory Group (OTAG) [R6] which is composed of representatives from the operation and users communities, from the middleware developers and from the JRA1 activity. The OTAG is the main supervisory body for the development progress. The prioritization step is particularly important in order to have a proper schedule of the development work.

Another project group, the User Service Advisory Group (USAG) [R7], focuses on the requirements for the complete set of user services run by the project among which are some tools (or parts of tools) developed by JRA1. Examples of these services are the VO Identity Card tool that is part of the Operations Portal and the GGUS Helpdesk system that is used by a variety of actors. Possible conflicts can happen between the decisions taken by USAG and OTAG on these “shared” tools, this eventuality was discussed during a USAG meeting and it was agreed that for all the tools but GGUS, prioritization can happen within the OTAG (where in any case user communities have representatives), while for GGUS dedicated USAG/OTAG joint meetings are needed. It was agreed that conflicts which cannot be resolved during these meetings be escalated to higher project levels; this did not happen during the first year of the project and hopefully will not happen in the future.

As of April 2011 about 80 requirements have been handled (and in open state) in the “Requirements queue” for the operational tools; five OTAG meetings [R8] were organized for their discussion and prioritization. About 15 requirements were already addressed and implemented in tool releases. Project metric MJRA1.4 measures the number of OTAG-approved requirements, while Project metric MJRA1.5 measures the average time for an approved feature to be released into production. MJRA1.4 is close to 100% for all the quarters since very few requirements have been rejected up to now, while the last calculated MJRA1.5 was 1.4 months.

Together with the “Requirement queue” a so called “Roadmap queue” was created in the RT system. This is the single access point to information about new releases for all the tools. Each new release is associated to an RT ticket and in the ticket the estimated release date, links to release notes, to



change-log and to documentation pages are provided. Requirements addressed by the releases are also referenced in the ticket as links to the tickets in the “Requirements queue”.



3 DEVELOPMENT AND MAINTENANCE

The development and maintenance work for the operational tools is done under the TJRA1.2. In this section the work performed and the releases produced by each Product Team (PT) during the first year of the project is presented with a brief description of the releases content; references are provided to documentation and release notes.

3.1 Operations Portal

Operations Portal releases were produced during all the four quarters of the PY1². In June 2010 the first Regional Package was released, it is an Operations Portal deployable on NGI instances and synchronized with the central one (section 4.1.1 provides more information on the Operations Portal regionalisation). During PQ1 porting to the Symfony web framework [R9] of some modules of the portal was started. This process resulted in two central instances running in parallel, the old one (inherited from EGEE and also referred as CIC Portal) and the new one based on Symfony. A detailed development roadmap for the tool (MS701 [R10]) was prepared in PQ1. A new release (V2.3) for the Central Operations Portal was produced during PQ2. The updates over the previous version consisted mainly of migrating the VO ID card module into the new Portal. The VO ID Card system records the life cycle of a given VO and links the VO managers to the project management for operations. Some improvements to the Operation Dashboard and to the web pages look and feel were also available in this version. New features related to VO Management were proposed as a prototype to be validated by the EGI VO administrators and by the NA3 activity managers that produced a series of comments and new features requests for the module. These requests are still under discussion at the moment of writing but most of them will be included in the PQ4 release. The Operations Portal VO ID card module will become the official way to register and update the static information for a VO. For more information please refer to the Operations Portal documentation pages [R11].

The migration of the broadcast tool to the Symphony framework foreseen for PQ2 was postponed to the beginning of PQ4 because the development of the new VO ID card took longer than expected (due to the many interaction cycles needed with the user support community which is the main user of the tool). The development work needed for the notification system and for the Lavoisier web service programmatic interface [R9] [R40], started in PQ1, was completed in PQ3 together with some improvements to the VO ID card. This resulted in two Operations Portal releases during PQ3, v2.4 and v2.4.1 – more details can be found online [R12].

One important upgrade of the regional package was made during PQ3. The synchronization process has been reviewed and now the regional portals are the authoritative source of information. The central instance is synchronized by pushing information from the regional instances.

PQ4 saw the release of the broadcast tool in the new portal (v2.5) and a refactoring of some modules (v2.5.1), including the Operation Dashboard that now has a new Central Operator on Duty (COD) view enabled. A new version (v2.6) was released at the beginning of April 2011 including new features and improvements for the Operation Dashboard (COD view included), for the Broadcast tool (mailing list reorganization), for the VO ID Cards (new registration page, glossary added, new help pages) and for the Portal homepage [R12]. This last PQ4 release also included the integration of external tools in the Operations Portal such as Bazaar [R55] and the YaimVO Configurator³.

² In this document PY is “Project Year” while PQ is “Project Quarter”.

³ A tool that lets you manage the last part of the YAIM configuration file – YAIM is the gLite configuration utility.



During OTAG-05 in December 2010 the integration of the Security Dashboard into the Operations Portal was approved. This dashboard contains security test results from external tools used by the the EGI Computer Security and Incident Response Team (EGI-CSIRT) [R52] such as Pakiti [R13]. Development for this integration started in PQ3 and continued in PQ4 to:

- Propose an adapted display and workflow to open tickets against sites in the Operations Dashboard
- Define/adapt/implement/ the XML (CSV,...) format of the reports for Nagios and Pakiti and make them available for the Dashboard
- Correlate and consolidate Pakiti and Nagios information
- Define and implement the mechanism of passing this information to the Dashboard
- Extend the Dashboard with the capability of displaying the information in the site view
- Establish the work-flow of the ticket creation (templates, assignment, RTIR ticket system integration)
- Integrate an additional Access Control List (ACL) model: Make sure proper authorization is applied (based on GOCDDB and EGI SSO) and make sure that EGI CSIRT/operations people can access all the data collected.

This is a major OTAG-approved requirement with an important impact on the initial development roadmap. This impact is further analyzed in Section 6.

3.2 EGI Helpdesk (GGUS)

Maintenance and development of new releases was performed during the entire PY1. Three minor GGUS releases were performed during PQ1, one per month and release notes are available at [R14].

The main achievements of these releases were:

- Regional view (xGUS) in production for NGI-DE since July 2010. This was a prototype to be evaluated by other NGIs – more details on the helpdesk system regionalization is in section 4.1.2.
- GGUS redesign: With the end of EGEE and the start of EGI-InSPIRE the GGUS website got a new logo and a new style sheet. Along with the trend of decreasing the height/width ratio of modern monitors (less height, more width) the navigation bar was moved from the top to the left of each web page. That way more space for the content is available.

Two GGUS releases were performed during PQ2 [R15], released respectively at the end of September and at the end of October 2010. These releases included:

- The integration of new NGIs and new VOs into the system
- The renaming and restructuring of various support units (reorganizing support units to fit the EGI model, adapt or remove legacy support units from EGEE)
- Introduction of new 3rd level support units (including IGE)
- Minor bug fixing
- New EGI logos and a different distribution of the web interface elements on the screen.

The xGUS portal was demonstrated at the EGI Technical Forum in September 2010.

During PQ3 the GGUS Product Team worked on the Deployed Middleware Support Unit (DMSU) workflow and on the first GGUS release containing the instances of the Initiative for Globus in Europe project (IGE) and EMI support units that are hidden from normal users and for which assignment is possible only by DMSU. On the xGUS side four xGUS instances were put online: one in Germany (DE),



Switzerland (CH), Serbia (AEGIS) and EUMEDGrid (external). NGI integration continued during the third quarter:

- 29 NGIs interfaced to GGUS
- 20 as support unit
- 5 with local ticket system
- 4 with xGUS.

A GGUS release is foreseen for PQ4 containing a second version of the middleware instance used for the third level support and a redesign of the Report Generator (a tool that create GGUS tickets statistics [R57]) according to client requirements and allowing for:

- More flexibility with new metrics available (i.e average and median of Response Time, Solution Time and Assignment Time)
- Output in further processable formats (for example xml).

Work on continuous integration of NGIs was also performed during PQ4 to get all NGIs integrated and to increase the number of automatic interfaces (local ticket system or xGUS).

3.3 Grid Configuration repository (GOCDB)

The release of GOCDB4 and decommissioning of GOCDB3 was the major goal of the GOCDB PT at the beginning of the project as a continuation of the work performed during EGEE-III. By design GOCDB4 provides:

- A GOCDB module (hereafter called “GOCDB4 Regional Module”) that can be installed by NGIs to store their own topology information
- A central system (hereafter called “GOCDB4 Central Instance”) to present all data collected from NGI instances
- A central input module (hereafter called “GOCDB4 Input System”) to insert data into the central instance. Data insertion into the central instance by design is possible also through the Regional Module but the synchronization mechanism between regional and central instances is still to be developed, more details in section 4.4.

Releasing and rolling out the GOCDB4 Central Instance and Input System was thus the first step towards creating a fully regionalized GOCDB since the GOCDB4 Regional Module cannot work by design (due to database schema changes) with the GOCDB3 Central Instance.

The GOCDB4 Central Instance was released in November 2009 in read-only mode and synchronized every ten minutes with the GOCDB3 read/write system. Given the high number of tools and middleware components depending on GOCDB information, the GOCDB3 decommissioning had to be carried out with extreme care. A release plan made of two major steps was defined:

1. Swap all GOCDB Programmatic Interface (GOCDBPI) URLs to point to GOCDB4 only
2. Release the frontend GOCDB4 input web portal.

Further details about the GOCDB4 rollout plan can be found at [R16].

During PQ1 the first step was completed after a large-scale production test which verified that all possible compatibility issues were solved [R17].

Server overload problems were found during the first hours of the test and were promptly solved by the developers. No further issues were found in the following days and the large-scale test was successfully completed.



The second step was completed in PQ2 on the 14th of October after an intense testing phase that involved many actors within and outside the JRA1 activity.

User feedback coming from various actors on this new release was collected on a wiki page [R18], minor bugs and cosmetic improvements were fixed on the fly when possible, otherwise bugs were opened to track the issues. However no blocking problems were found. More information, technical details and the outcome of this process are available at [R19].

During PQ2 and PQ3 effort was spent on the following:

- Providing a production quality packaging of the regional module
- Data access optimization
- Providing a SOAP [R56] interface in parallel with the GOCDBPI.

In November 2010 the main GOCDB developer left JRA1 and during PQ3 the new developer had to familiarize himself with the code and with the Oracle PL/SQL language used for the tool. Patches were added to the production servers to support secure SSLRenegotiation (required to support new browsers such as Firefox4 and Chrome8). Missing database columns were added. Lots of bugs were fixed, mostly regarding regular expression and validation rules.

GOCDB status and requirements were presented at the OTAG-06 meeting [R20] where key requirements were prioritized and different options for their implementation were drafted. These were discussed in numerous phone conferences (both internal and external to JRA1) to verify the impact on other tools and on the middleware. Most of the effort in PQ4 was spent to address these requirements at design and implementation level. Examples of these key requirements are:

- Record Certification Status Histories (who, when, audit table) with PI/GUI updates.
- Extensions to query historic (decommissioned) NGI-to-Site associations.
- New kind of downtime status in GOCDB for Early Adopter sites (EA, see [R3][R4]) for adapted reliability metric calculations.
- Clean roles in GOCDB and added Chief Operation Officer role
- Mask sites from different/entire communities (related to regionalisation).
- Naming schema compliant with GLUE2.0 specification
- Naming schema able to encompass UNICORE services (more details in section 3.10).

As for the other JRA1 tools it was decided to move the GOCDB requests tracking tool from CERN Savannah to the EGI RT system.

3.4 Accounting Repository

During PQ1 focus was on the integration of the APEL accounting system [R21] with the message broker network as decided during the previous EGEEIII project. The ActiveMQ based APEL server has been consolidated to a production level to accept and process records through the newly released glite-APEL client.

During PQ2 the APEL tests have been successfully migrated from the previous legacy submitting system to the SAM Nagios and the implementation of a summary records consumer with a clear defined message format started in PQ1 was almost completed. An update with bug fixes for the gLite-APEL service has been certified and released into production. Work on the design of a distributable Regional Accounting Server has started and during PQ3 the design of the regional APEL system has necessitated several redesigns:

1. **The Schema:** has been simplified by removing unused or duplicated fields; new fields added for MPI; fields renamed to align with Open Grid Forum's Usage Record [R47][R54]; new record



primary key defined for integrity, and a new Job Summary Record defined based on the proposed extension.

2. **Message Format and Infrastructure:** change to use the STOMP protocol [R53] on the EGI production messaging infrastructure. This required a new encryption and authorization model (x.509 based).
3. **The job record database** was changed to reflect changes to the schema and message format.

During PQ3 the design of the above was done and the development started.

A roadmap has been defined to roll out a new infrastructure in PQ4. This will start with a test infrastructure to enable testing of a new gLite-APEL in EMI and clients developed by partners in other grids who used to publish by direct database insertion and will now publish Job Summaries using ActiveMQ. By the end of PQ4 the new infrastructure will be in production and a gLite-APEL using the new features described above will have been released by EMI, and hopefully EGI.

3.5 Accounting Portal

Problems with filling vacancies at CESGA due to administrative and national law issues delayed the planned release dates for the Accounting Portal. The hiring process was completed at the end of PQ3 when two people were contracted for the work.

Nonetheless, work was performed during the first two quarters to support the new Programmatic Interface released with GOCD4. Tier2 report pages were updated and several minor issues solved, the NGI View was added to the central accounting instance and an improved installation support was added to the regional package (more information on the regionalised Accounting Portal is given in section 4.1.6).

No official releases of the accounting portal were performed during PQ3, but development effort was spent to migrate the VO information: switch from direct Oracle connections to XML feeds in the Operations Portal and modification of the accounting scripts to obtain the data correctly. In PQ4 the Accounting Portal requirements stored in the EGI RT system that came from various clients were reviewed in order to create a development work plan at the beginning of PY2 in collaboration with the OTAG.

3.6 Service Availability Monitoring

During PQ1 of the project three releases of the SAM system were produced: SAM-Update 1/2/3. Details and release notes can be found in the CERN JIRA task tracker [R22].

The transition from the ROC model (SAM instances deployed in each ROC) to the NGI model (SAM instances deployed in each NGI) was almost completed in the first quarter (see Sections 4 and 4.1.7 about SAM regionalization).

The procedure to validate an NGI SAM-Nagios box using a ROC level NAGIOS was created [R23] to be exploited by the SA1 activity.

Starting from SAM-Update2 the tool was released using the EGI official release procedure [R3][R4] with one Early Adopter site (a production site that pre-deploys the service for a specified period giving the green light for full deployment).

Two updates were released in PQ2, a minor one (update-04) at the beginning of September 2010 and a bigger one at the end of October 2010 (formally containing two updates, update-05 and update-06).

Update-04 major achievements were:

- Metric Resource Store schema updates



- First version of MyEGI bundled, but not activated
- Improvements in debug and signal handling in probes
- APEL test integration
- Robustness improvements in msg-to-handler.

Update-05 and Update-06 major achievements:

- Merge databases of different components to a single DB
- Use the ATP component [R24] for Topology including features for VOfeeds [R25]
- First (activated) release of MyEGI (in parallel with myEGEE for a while).

Besides the development done to release these important updates, effort was spent in PQ2 in helping middleware providers to develop probes for new middleware types and to start their integration into the SAM framework. Details on this activity are given in Section 3.10.

During PQ2 development started that was completed in PQ3, of a new probe to check Certification Authority validity that does not need to be updated on every CA update. Discussion with the SA1 activity on how to aggregate physical sites and services into virtual sites for monitoring purposes started during PQ2. This discussion resulted in a RT requirement discussed and approved at the OTAG-04 [R20] meeting. This requirement heavily impacts the ATP component development and its implementation and is still ongoing.

There were three updates in PQ3, two public and one internal. Update-07 was released at the end of November 2010. Update-08 was marked internal, while Update-09 was released in February 2011.

Update-07 major achievements were:

- Use of ATP as topology provider for the Nagios Configuration Generator (NCG) component
- First release of new CA distribution probe
- Integration of ARC probes.

Update-08 and Update-09 major achievements:

- MyEGI standalone central instance
- MyEGI web services
- Added SAM release to information provider
- Support for robot certificates
- Support for configuration of uncertified sites
- Yaim (the configuration system for the tool) cleanup
- Support of MPI services in topology.

A new update (SAM-Update10) entered the official release procedure on April 11th and, at the time of writing, is held in the staged rollout phase waiting for a bug discovered by the Early Adopter site to be fixed. This update contains various improvements and the support for new VOfeeds that are the starting point to address the Virtual Sites requirement previously discussed.

For more information on these SAM updates, release notes are available at [R26].

Integration of ARC probes (Update-07) required additional testing and documentation was provided to SAM administrators [R27].



As already mentioned the SAM team within JRA1 is no longer responsible for the middleware probes development, but only for their integration into the monitoring framework. The handover of this process was presented and discussed at the EMI All Hands meeting in Prague. EMI accepted the handover of Nagios probes relevant for products provided by EMI. At the end of January it was agreed to create an EGI-EMI task force that would start the handover of probes. The responsibility of the EGI side is to identify which probes are needed and provide pointers to specifications and current implementation.

Additional support was provided to NGIs that started deploying VO SAM instances (IBERGRID and France). Based on requests the new VO configuration profile was added to Update-09.

3.7 Metrics Portal

Being developed at CESGA, the Metrics Portal suffered the same problems related to the hiring process that affected the Accounting Portal for the first three quarters of the project and no releases were performed during that period. Minor fixes were done to the portal inherited by EGEE-III to accommodate changes in underlying data source tools such as Gstat [R35]. Input from the project was needed on how to evolve the tool to encompass the new project structure and performance metrics. Once the staff hiring problems were resolved the input was collected and a detailed 3 year development road-map with schedule for the new EGI Metrics Portal was created [R28] and presented at the OTAG face-to-face meeting in January 2011 [R20] and to the AMB face to face meeting in March [R29]. The presented development is based on a spiral model: three complete cycles are defined, producing a new release of the metrics portal at the end of each cycle. Inside each cycle there are six steps, one devoted to each data source. After each step a preview release including the additional metrics created will be available in the development version of the metrics portal for internal review. The first preview release was released at the beginning of April 2011.

3.8 Broker network configuration and infrastructure

During PQ1 and PQ2 the main effort was on the implementation of the requirements from the APEL team for message brokers. This required the inclusion of the authorization plugin on the brokers on which resulted in the denial of all broker-to-broker communications. Broker-Update-2 was released with two new features to meet the APEL broker requirements of EGI:

1. Addition of GOCDDB support to activemq-voms2users package
2. Addition of group support to the activemq-voms2users package.

Documentation was produced on the current production message broker network's authorization needs and it was proposed the usage of authorized only topics [R30].

On the infrastructure side continuous effort was spent to maintain the development repository and the build system for SAM and to investigate the possibility of integrating the build system with the JIRA used by the SAM team in order to avoid delays on the new package builds. This was unsuccessful due to the lack of information provided by the JIRA API (i.e. the responsible person per package or the JIRA account <-> x509 subject DN mapping) so this was dropped.

3.9 Ops Tools migration to egi.eu domain

During the first year of the project all the JRA1 tools underwent a rebranding process to add EGI logos and copyright in the web interfaces. The rebranding included moving the central instances URLs to the egi.eu domain. Quick links to all the central and regional instances are available at [R31].



The decommissioning of the old gridops.org domain was decided during PQ3. The decommissioning date was initially set for mid March 2011 and then postponed to mid June 2011 to wait for a new gLite-BDII release that solves potentially disruptive dependency on that domain.

3.10 Integration of new middleware types

One of the main objectives of the EGI-InSPIRE project is to include new middleware types into the production infrastructure. This has of course implications on the tools, in particular on the monitoring system (new probes are needed) and on the GOCDDB (new service types and capabilities have to be taken into account). Since the beginning of the project the support to the development of probes for new middleware types and their integration into the SAM framework has been a high priority task for JRA1. SAM-Update7 [R26] in November was the first update to include in production probes for the ARC middleware. UNICORE probes development and integration is still ongoing at the time of writing but in an advanced state and will be ready in one of the next SAM Updates. GLOBUS5 probes are already included in SAM and only need to be tested once a Globus site will join the infrastructure. JRA1 also had preliminary contacts with the European Desktop Grid Infrastructure project to trigger the same activity for their middleware.

Impact on GOCDDB is particularly important for the UNICORE middleware mostly because of the need to store dynamical information such as services endpoints that are published by the production infrastructure information system. This information should be stored on a dedicated registry developed within EMI, but this is not available yet and GOCDDB is intended to store static topology information only, so a medium term solution has to be found. This was preliminary discussed at the OTAG-06 meeting where a couple of solutions were proposed by JRA1. The final decision was committed to a dedicated task force on UNICORE integration that is still working and holding regular meetings: their outcomes are updated on a dedicated section of the EGI wiki [R32].

3.11 Monitoring tools availability

A development activity that involved all JRA1 product teams was established since the beginning of the project to produce SAM probes to monitor the operational tools themselves (i.e. web pages are responding and basic functionalities working). These probes will be initially integrated on a SAM server at CERN [R33] to ease the work of tools administrators and to monitor the tools availability and reliability. The activity is tracked through RT tickets in the JRA1 queue and basic probes that check html page responsiveness are now available for all the tools, while more complex probes are still under development. It is important to have availability and reliability numbers for operational tools also to understand if their current failover systems need to be strengthened.

3.12 Dependencies and interaction with externally developed tools

The EGI-InSPIRE projects relies, even for key activities, on operational tools developed outside JRA1 and outside the project itself. Examples of these tools are the Availability Calculation Engine (ACE) [R34] - used to produce availability and reliability numbers for the production sites, Gstat [R35] – a visualization and aggregator tool for the Grid information system, GridView [R36] – a historical visualization tool for the SAM system results. There are a series of cross dependencies between these tools and those developed within JRA1, for instance all the “external” tools depend on the GOCDDB as topology information provider while the JRA1 Metrics Portal depends on Gstat to get aggregated statistics from the information system. Moreover, the EGI-InSPIRE project has development requirements, enhancement and feature requests for external tools, in particular for the ACE system, but since these tools are being developed outside the project, EGI has no control on



their development roadmap and thus it is of utmost importance to have a discussion forum to discuss issues and requests from and to both internal and external tools.

During this first year of the project the situation was progressed through adhoc phone conferences and private communication when needed, but the plan is to have a dedicated, cross projects body for this purpose. This body could be the old EGEE Operation Automation Team (OAT) with a new mandate.

The definition of formal MoUs between EGI and the other projects developing operational tools (i.e. WLCG) could also help in this respect.



4 REGIONAL MODEL FOR OPERATIONAL TOOLS

4.1 Background

EGI attempts to have a strong regionalisation model in the area of operational tools. Significant effort was devoted to this activity during the second year of the EGEE-III project, providing a regionalized instance of each of the tools or a regionalized view on the central instances.

At the end of EGEE-III the SAM Nagios-based monitoring was the only tool fully regionalised. This meant that a Nagios instance, including a MyEGEE portal, was active in each region. To complete the transition at the beginning of the EGI-InSPIRE project the former ROC instances needed to be decommissioned in favour of NGI instances in order to reflect the NGI-based operation model of the project.

During EGEE-III regionalisation plans were created for other tools such as the Operations Portal and the GOCDB. The development to address these plans was started, but its finalization, testing and release to production needs to be completed within the EGI-InSPIRE project.

TJRA1.3 is the JRA1 task responsible to complete the regionalisation of the operational tools and it is a one-year only (PY1) task.

At the beginning of the project the NGIs' deployment plans for regional operation tools were collected in collaboration with the SA1 activity and the resulting matrix is available on the EGI wiki [R37]. Table1 summarizes the NGIs' answers:

	Wish to install at NGI level	Wish to keep it centrally	Not decided	NGI has custom solution	N/A
Accounting portal	10	15	5	2	2
GGUS (xGUS)	4	15	1	13	1
GOCDB	11	21	0	0	2
Operations portal	10	21	0	1	2

Table 1 - Summary of the NGIs plans for regional operational tools deployment (N/A means number of NGI that did not provide feedback for the tool).

The main objection to the current implementation of regional tools is that they are not easily customizable by NGIs and are just copies of the central instances deployed locally. To face this issue (among others) during the Operation Management Board and OTAG face-to-face meetings in January 2011 it was decided to create a joint SA1/JRA1 Regionalisation task force in order to better define use cases for regional instances of operational tools. The outcome of this task force will be mirrored in clear requirements to drive JRA1 regional tools development.

4.2 Operations Portal and Dashboard

The development of a portal regional package started in EGEE-III: the goal of this work is to distribute the different features and modules of the Central Operations Portal into a regional package for NGIs to cope with NGI needs (i.e. integration with other regional tools).

The first release of the first production package was on June 8th 2010 and it included the first release of the dashboard module. The different parts of the package and the documentation are distributed via a SVN repository.



Currently four instances of the regional portal are installed into production for NGI_IBERGRID, NGI_GRNET, NGI_CZ and NGI_BY.

Details on the regional architecture and synchronization methods between regional and central instances can be found in project milestone MS703 [R38].

This regional package is currently using the same central data sources as the central instance, but is designed to integrate the regional version of GOCDDB and regional helpdesks if they are deployed. This regional package will be constantly improved during the coming year to add some of the other features that will be integrated step by step into the Central Instance. These features will be added one by one into the central instance, and if relevant to a regional instance an update for the regional package will be provided as well.

4.3 EGI Helpdesk

The GGUS portal is the central helpdesk of the EGI Grid. It acts as an integration platform between other regional helpdesk systems that are in production. The set-up with a central tool and various different integrated remote systems has been in place since the beginning of EGEE. All project-wide support units are present in the GGUS system; regional or specialized topical support units can be reached via GGUS through the interfaces their respective ticket systems have with GGUS.

One of the major changes to the user support infrastructure caused by the transition from EGEE to EGI is the move from ROCs to NGIs. This means that instead of 12 regional ticket systems (one in each ROC) now every NGI has to set up the respective user support infrastructure, and interface it to the central GGUS system. This of course increases the number of interfaces that have to be maintained in GGUS. In order to cope with this it is necessary to strictly standardize the interfaces to new NGI helpdesk systems. Only web services and messaging interfaces are being accepted now. The interface standards for these two methods are described in the documentation collection of the GGUS portal [R39].

In EGI we have identified three options for how NGIs can connect to the central user support infrastructure:

- NGIs can set up a regional ticket system and interface GGUS via web services or messaging. This option is currently used by five NGIs.
- NGIs can use the GGUS system directly. In this case there is only one support unit in GGUS for the whole NGI. This option is currently used by 11 NGIs.
- NGIs can use a customizable version of GGUS called xGUS. This was developed to give NGIs a simple way to set up a ticket system without having to host it themselves.

The design phase of a customizable version of GGUS started in 2009 during EGEE-III, resulting in resources being assigned to developing such a system in the proposal for EGI-InSPIRE under JRA1.

The xGUS system has the same basic functionality of GGUS and is based on the same workflow engine. xGUS consists of a web front end that can be customized with a NGI-specific banner and a backend that includes all basic support portal functionalities such as a ticket database, an email engine, user administration, a news module and basic portal administration. This means that fields like 'Type Of Problem', 'Affected Site', 'Affected VO', 'Responsible Unit' and link lists can be administered directly via the portal.

The template comes with synchronization to GGUS, so tickets coming from or going to GGUS are duplicated and synchronized automatically.

xGUS is hosted centrally by the GGUS team, relieving the NGIs from having to maintain a ticket system locally, so it is not meant to be deployed locally. The first instance of xGUS was set in production for the NGI-DE, the German NGI, and the evaluation from other NGIs started during PY1.



The xGUS system undergoes the same release procedures as the GGUS system, as the two systems are closely related.

4.4 Grid Configuration repository (GOCDB)

The last released version of the GOCDB regional module dates from October 2010 (gocdb4.0d7). It is a fully functional tool, and all the basic components are included so that initial deployment can be tested using that version. The documentation is available on the EGI wiki [R41]. GOCDB regional modules can be operated in full production, but can not yet synchronize to the central GOCDB. NGI (ROC) plans about regional module deployment were investigated back in 2009 and three kinds of regions can be identified:

- Region 1: deploying and using a distributed, customizable version of the current GOCDB
- Region 2: continuing using the central GOCDB
- Region 3: deploying/using another tool than GOCDB, interfaced with central GOCDB for interoperability purposes

The regionalisation model for the GOCDB is currently a controversial topic; possible use cases and implementation options are still being discussed within the regionalisation task force already quoted in the introduction of Section 4. Future GOCDB regional module development will be focused in the following directions:

- 1) Data Grouping/Scoping:** to separate EGI and non-EGI data/sites - A new group named EGI will be created along with tools to insert/remove data to and from a specified group. With this functionality in place NGIs will be able to store and retrieve both local and EGI level sites in the central GOCDB. This will allow them to centrally include/exclude sites/services from the whole EGI infrastructure as required (provides extensible data visibility/scoping according to groupings) so that regional GOCDB installations are not required (e.g. for NGIs without effort to install/maintain a regional GOCDB).
- 2) Regional-Standalone GOCDB:** to release a standalone instance of the GOCDB for regional NGI use. This regional instance will be supported regardless of other developments. These instances will be customisable by NGIs willing to do so.
- 3) Regional-Synchronised GOCDB:** this is a longer term development. A regional-standalone GOCDB instance will synchronise its EGI-scoped regional data with the central instance. This would allow NGIs to deploy their own GOCDB, customise it and publish EGI data (constrained by the XML schema) to the central portal. Publishing this data would have to be transactional, most probably via a schema constrained WS/REST interface (preferable over asynch messaging which would require response queues). The central portal would be the definitive source for all EGI information. It may be necessary to have two-way synchronisation, where the central instance would also update regional instances since NGI extensions and custom data would not be synchronised with the central GOCDB, therefore point 1) is necessary to complete the synchronisation.

Since the regional model is still under discussion a page in the GOCDB wiki is dedicated to this topic [R42] and updated information on regional GOCDB development can be found at [R43].

4.5 Accounting Repository

The APEL regionalized architecture was presented as a poster in EGEE09 conference and is available from the APEL home wiki page [R21]. This architecture is still being followed. The three types of

region were identified (an NGI or group of NGIs is a region), region A (direct publishing) is now in place (see Figure 3 below). Work has started on the other two types with the development of a summary publishing client (required for both type B (distributed APEL) and C (own accounting solution, APEL regional publisher)) the development of summary ingestion in the server and a redesign of the repository to make it more suitable for distribution.

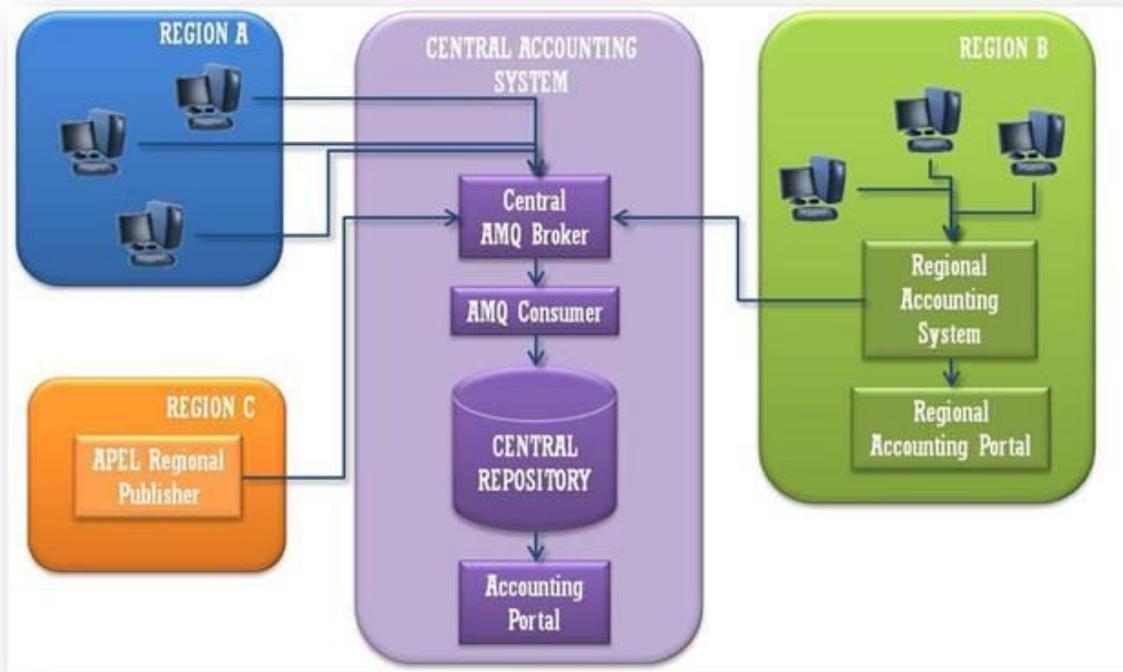


Figure 1 -Accounting Regionalisation

The schedule for the development and release to production of the regional APEL system are summarized Table2:

1	Complete migration from R-GMA to ActiveMQ	31-Jan-11	SA1
2	Design new DB schema and Message Format for jobs and summary	21-jan-11	JRA1+EMI
3	Pre-production version of summary infrastructure	11-feb-11	JRA1
4	Summary Infrastructure deployed at RAL	25-feb-11	JRA1
5	New Central Infrastructure in Production	02-apr-11	SA1
6	Central Server ready to receive new format input	15-Apr-11	SA1
7	New APEL Publisher (stomp-based) released	16-Apr-11	SA1+EMI
8	Regional Accounting Server packaged and released to NGIs.	01-Jun-11	SA1
9	Regional Deployment and Migration	June onwards	SA1

Table 2 - Development schedule for the regional APEL system



4.6 Accounting Portal

The prototype version of the regional accounting portal is already available. Currently several NGIs have expressed their interest in deploying a regional instance of the regional portal.

During the following months the installation scripts will be improved making it easier to deploy the portal. Additional improvements will be incorporated based on the feedback from the early adopter NGIs. New functionalities incorporated into the central accounting portal will be ported to the regional accounting portal.

The regional accounting portal will use the new accounting infrastructure based on Messaging. This new accounting infrastructure allows greater flexibility to configure regional accounting repositories and several modifications are foreseen in the regional accounting portal to support ActiveMQ. Besides the regional accounting portal, a country view is offered in the central accounting portal that allows NGIs to review the usage in their region. This view will be continuously improved based on NGI feedback.

4.7 Service Availability Monitoring

The monitoring infrastructure based on Nagios and Messaging is already fully regionalized as discussed in the introduction of Section 4.

4.8 Metrics Portal

Currently there are no regionalization plans for the Metrics Portal.



5 ACCOUNTING FOR DIFFERENT RESOURCE TYPES

The evolution of the accounting system to encompass the new different resource types that will join the production infrastructure and to support an economic model needed for the self sustainability of the EGI Grid is an objective of the JRA1 activity. The task responsible for its development is TJRA1.4 – Accounting for different resource types – which is a 3 year task starting from PY2, so no official development was carried on during this first year. The new accounting types will include:

- Application usage: account for the use of particular applications by specific communities which would help resource providers, VO managers and application developers to improve the services they offer
- Data usage and storage: account how much data is accessed by different communities and how much space is occupied by data belonging to those communities
- Accounting of high performance computing and cloud computing usage.

No official development was done during the first year, but prototypes for storage resources have already been implemented by the LUH partner.

Preliminary discussions were initiated at the accounting workshop held at the EGI Technical Forum in Amsterdam and will be elaborated in greater detail in dedicated meetings that will be organized as the new project year will start.



6 INTEGRATED OPERATIONS PORTAL

A task (TJRA1.5) within the JRA1 activity is completely dedicated to the evolution of the Operations Portal towards a service-oriented model by providing services related to all types of information handled within the Portal. The initial plan covered the porting to the Symfony [R9] open-source web-development framework, the extension to support messaging (downtime and broadcast tools) and the harmonization with other portal frameworks, the modularisation in order to make its views pluggable through portlet/widget technologies into other portals (i.e. VO portals) and mobile devices. The following sections report a summary of the development and design work done within this task during the PY1 on various components of the operation portal.

6.1 *Porting to the Symfony open-source web-development framework*

This phase is almost achieved and all services will be provided in the Operations Portal Framework by the end of the first year of the project. The Symfony solution has been chosen for maintainability and development efficiency reasons, but also for the modularity given by the framework. This modularity will allow for an independent functioning of the different modules even in a service-oriented model (a service will be equivalent to a module).

6.2 *Towards a service oriented model*

To allow for a service oriented model the architecture of the portal has been built in order to easily extend the number of data sources, and to propose a standard access to information. The core of the data gathering system is a web service facility called Lavoisier [R40]. In the case of known technologies, new views will be added by using an existing plug-in out of the wide range already available, while for new providers new plug-ins will be developed to be able to retrieve information. Lavoisier's flexibility allows the Operations Portal to be ready to integrate almost any kind of new information if needed.

Moreover, a programmatic interface based on the Web service of Lavoisier has been developed. The idea is to provide the whole set of information provided by the Portal in standard formats. The current interface provides information related to NGI and VO in an XML format. This solution will be extended to support:

- XPath, XSLT, XQuery and STX query languages
- XML, JSON and YAML formats for query results
- SOAP, REST communication protocols
- STOMP, OpenWire messaging protocols.

To achieve these goals a new major release of the Lavoisier system (v2.0) is foreseen to be completed in June 2011. The integration step will last from July to December and the PI will be enhanced step by step with the use of new possibilities offered in this version.

6.3 *Interoperability with GOCDB*

Another important goal of the TJRA1.5 task is the integration of the Operations Portal and GOCDB under a common interoperable toolkit for Grid operations, which can be achieved in two different ways:

- Integration of a common central human interface allowing users to access both central services through a single entry point.
- Integration of interoperable back ends for distribution to NGIs as a single package if possible.



Such development will require effort at the data representation level as well as at the interface and data transfer level. Both Operation Portal and GOCDDB development teams aim at providing a common web interface to handle information coming from both tools so that the underlying messaging is transparent to the user. A study on how the harmonization of the two tools can be accomplished started during this first project year. The preferred vision of the development teams is to have “independent tools working together” rather than “one single tool”, keeping the integration transparent to end users that will not care about how this is done in the backend. To the end users, the solution will be seen as “one single tool”.

The following guidelines should drive the implementation:

- Keep modularity – avoiding to create a bulk monolithic tool
- Allow one part of the system to be installed and operated without the other
- Reuse the legacy wherever possible (no time, no money to redo everything):
 - The tools use different backends: Oracle vs MySQL
 - DB schema and principles are different (use of PROM for GOCDDB).

Two solutions were considered in the implementation study:

- a) **Merged front-ends and separated but interconnected back-ends.** Back-ends will interact with one another to join information about services and VOs. This solution has the benefit to reuse legacy databases while still providing a single top-level interface and maintaining great modularity: the deployment of both tools will not be required. A drawback is the complex failover solutions that are needed to achieve a fault tolerant, high availability tool.
- b) **Fully merged front-ends and back-ends.** This solution implies a complete redesign of the database model integrating data currently handled by both GOCDDB and CICDB (Figure 2) and the development of a new interface to this database as a single front-end. Benefits are a less complex failover strategy and an easier implementation of cross queries. Drawbacks of this solution are that the resulting tool will be more complex and hence more difficult to maintain, with less modularity. It also implies a complete redesign of the legacy implementation.

The currently implemented relationship between GOCDDB and Operations Portal components is shown in Figure 2 while the proposed solution a) is depicted in Figure 3. At the time of writing solution B seems to be very complicated and would need a great deal of work probably beyond the possibility of the GOCDDB and Operations Portal product teams given the effort that they currently have in JRA1. A discussion with the Advisory Group is needed in order to take a final decision – this will be done in one of the next OTAG meetings. This work will start in January 2012.

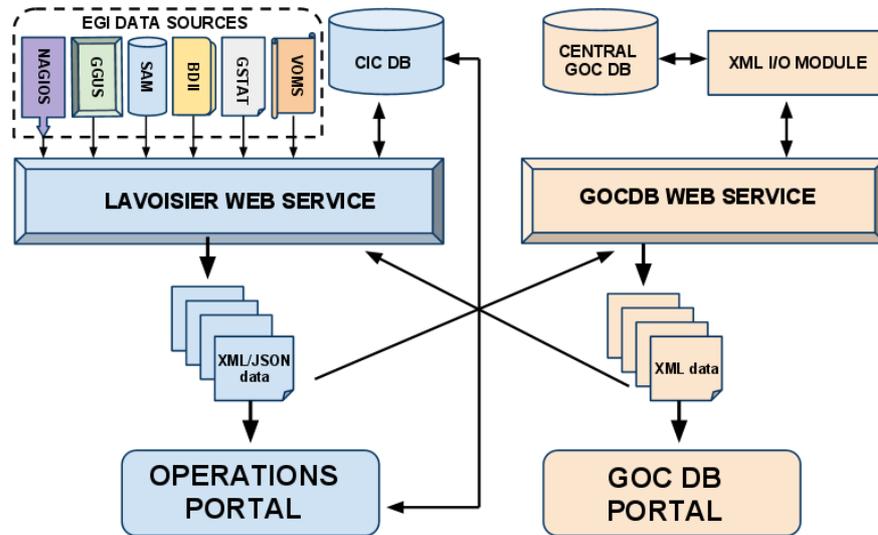


Figure2 – Current relationship between Operations Portal and GOCDB components

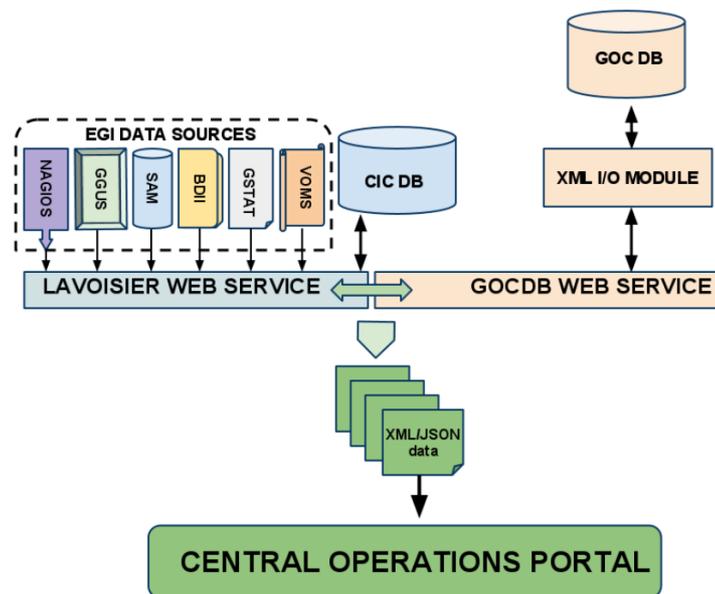


Figure 3 –Operations Portal and GOCDB harmonization envisaged solution. The solution envisaged has RAL hosting the GOCDB database and the Operations Portal in Lyon providing interfaces in the Symfony framework.

6.4 Changes to the initial workplan

More information on the integrated Operation Portal roadmap was given in project milestone MS701 [R10] which will be amended and updated annually. The initial release scheduled had to be modified during this first project year due to various reasons, among them:

- The validation/discussion about the VO ID card developments with the user support communities was longer than expected
- The time spent on the support for the Regional packages was underestimated



- New tasks and big requirements have been introduced during the year, that is, important changes on the COD view and the integration of the security dashboard (this will be clarified later in this section).

6.5 Security Dashboard

The development plan and requirements discussion for the Operation Portal took place during the OTAG-05 phone meeting. Its main outcome was the approval of the integration of the security dashboard used by EGI-CSIRT to detect security vulnerabilities at sites within the Portal. The goal of this integration is to detect and inform sites and ROD/COD managers about security incidents and vulnerabilities acting as a single access point to information for both operational and security issues integrating a workflow to handle open tickets against sites. The security dashboard is based on Pakiti [R13] its information will be correlated/aggregated with the data coming from SAM. An ACL model needs to be studied and implemented to make sure that proper authorization is applied (based on GOCDDB and EGI SSO) and to make sure that only EGI CSIRT people can access the sensitive collected data. At the moment of writing a first prototype for the integrated security dashboard is foreseen to be available at the end of May 2011.

6.6 COD View

The model of the central view has been adapted to central Operators needs.

- Sites are grouped by ROCs/NGIs
- COD can ticket against NGIs/RODs
- The templates and the different escalations steps have been changed.

This work has been done between February and April 2011 .



7 FIRST YEAR ISSUES AND MITIGATION

The following main issues affected the JRA1 activity during PY1:

1) **Staff issues in various development teams:**

- As mentioned previously there were hiring problems due to national law restrictions at CESGA (responsible for both Accounting and Metrics Portals). Development of new features and the redesign of the Metrics Portal could only effectively start during PQ3.
- The main GOODB developer left RAL at the end of October 2010. He was promptly substituted but during the first period the new developer inevitably had to familiarize himself with the code.
- The main myEGI developer left CERN at the end of February. He had effort only in TJRA1.3 that is a one year task (PY1), but a maintainer for myEGI needs to be identified from within the SAM team.

2) **Absence of 2nd level support units for the tools within the project.** Within the EGI project the second level support for middleware components is done by experienced teams of the Deployed Middleware Support Unit (DMSU), a dedicated unit funded by the project. The equivalent of the DMSU experts for operational tools were not foreseen by the project and the second level support is done by the developers together with the third level support, overloading the PTs of the fully regionalised tools (SAM team in particular). Actions have been taken in order to mitigate the situation. A DMSU support unit was created for this tool and volunteers were searched for within the operations community. Currently two people volunteered to participate to the activity joining the developers that are already working in the second and third level support. The sam-support<at>egi.eu mailing list was created for the support unit.

3) **TJRA1.3 ends with PY1** but the regionalisation activity is not completed and regionalisation use cases are still under discussion at the Regionalisation Task Force as presented in section 4. This discussion will reflect in potentially big requirements to the tools development but the matching effort ended in PY1. Together with the project office the JRA1 activity leader will investigate if all the TJRA1.3 effort was spent and if it can be redistributed. If all the effort was spent a backup plan is needed.

4) **Communication with other operational tools:** there are no official communication channels and discussion forums with externally developed tools such as GSTAT, ACE and GridView. The foreseen mitigation is to re-start the activity of the EGEE-III OAT group with the new mandate of being a forum for cross projects tools. Up to now the activity of this group was not started. MoUs between projects, especially for what concerns the ACE component, that is vital for the production infrastructure, could help in developing requirements for these tools.

5) **Authorization in messaging infrastructure:** There has been a slight concern over delays in the production messaging infrastructure supporting the authorization requirements of APEL. As mitigation at the EMI All Hands Meeting it was decided it was better to do authorisation on the APEL side based on the digital signing of the message with the client host certificate and then encrypt the message with the central private key so that when it is decrypted it still has the digital signature for auditing purposes.

6) **GGUS workflow:** GGUS has many different users coming from different projects and communities, i.e. EMI, IGE and WLCG. It is also the most exposed tool to end-user



communities so USAG is also dealing with GGUS. All the involved actors are making requirements to GGUS that now has different requirement workflows. The situation became difficult to handle by the product team, in particular in the prioritization part. As mitigation the OTAG-07 was called as a joint meeting with representatives from all involved actors to agree on a common workflow and discussion forum. The common workflow based on RT is currently under evaluation. Joint OTAG/USAG/WLCG meetings were agreed as a discussion forum for big requirements while weekly phone conferences between developers and interested actors will follow the day-by-day development and urgent issues.



8 ATTENDED/ORGANISED CONFERENCES AND MEETINGS

The JRA1 activity started with a face-to-face kick-off meeting organised in June 2010 in Amsterdam and continued direct communication was possible through periodic phone conferences [R44]. Another JRA1 face-to-face meeting was held in January 2011 in Amsterdam co-located with the OTAG-06 and the periodic OMB [R45] meeting. OTAG meetings are organised and chaired by JRA1 [R20].

Attendance on other project bodies and advisory groups (i.e. USAG [R7], OMB, UNICORE integration task force [R32], Regionalisation task force [R46]) was guaranteed by the activity.

JRA1 staff actively attended the EGI Technical forum in Amsterdam participating in various sessions and organising an operational tools dedicated panel [R51]. An accounting workshop was also organised at the technical forum by the TJRA1.4 task leader.

Five people from JRA1 actively attended the EMI all hands meeting in Prague in November 2010 exploiting that event to interact with developers on middleware/tools across topics such as SAM probes development.

Six JRA1 members attended the EGI User Forum in April in Vilnius, but no official JRA1 or tools dedicated meetings were organised within the event.

Open Grid Forum [R47] and International Symposium on Grid Computing [R48] events were also attended by JRA1 representatives in particular for accounting-related topics.



9 CONCLUSION

This document described the work performed by the EGI-InSPIRE JRA1 activity from its early stages up to the end of the first year of the project. The JRA1 activity is responsible for developing and maintaining some of the operational tools needed by EGI and is composed of development teams distributed across Europe mostly inherited from the previous EGEE projects, where a dedicated activity for operational tools development was not present and the teams were embedded inside the SA1 (operation) teams.

The set up of the activity was done following a model of centrally coordinated independent teams - independent in terms of development and testbeds, code repositories, bug tracking tools, build systems, documentation and distribution facilities. This model had the main advantage of minimizing the startup time of the activity, allowing the development teams to be immediately productive; thus JRA1 was able to release updated versions of the tools from the very first stages of the project. A drawback of this approach was the scattering of information regarding the tools that was mitigated by creating regularly updated and detailed pages on the EGI wiki and RT facilities. It was decided to create a common requirements workflow for all the tools. This process was agreed not only within the activity but also with other project activities that needed to handle requirements (i.e. NA3, SA1) in order to give uniformity to the process of gathering and accepting new feature requests from clients and users. This workflow is based on the EGI RT system and JRA1's development teams need to migrate towards this tool, but the common workflow is now in place. This document has described how the software development and maintenance work was carried out on all the tools during the first year with many new versions released during the year.

The main achievements within JRA1 were, among others: i) the release to production of the GOCDB4, which constituted an important step towards the regionalisation of the tool; ii) the refactoring of important SAM monitoring components with the release of the new myEGI visualization tool and the inclusion of ARC middleware probes; iii) a complete refactoring and regionalisation of the Operation Portal with the inclusion of the security dashboard; iv) the renaming and restructuring of GGUS support units and tickets workflow to accommodate the new EGI structure - the xGUS system was also released; v) the migration of the accounting system to the messaging infrastructure for the data transport layer; vi) the creation of a completely new metrics portal and vii) the migration of all the tools to the egi.eu domain and the rebranding of the web pages to show EGI logos.

The SAM tool, being fully regionalised and deployed in all the NGIs, was released following the official project release procedures, which are also applied to the third party middleware components. One of the main achievements during PY1 was the release of the regional package for the Operation Portal which is already deployed in four NGIs. However, regionalisation use cases and interaction models among the tools are still controversial and currently discussed in a dedicated task force which will provide requirements and guidelines to the regional tools development. Most of the effort for regionalisation was provided by an activity task which ends with PY1 and how to cope with this will be one of the major challenges of the next year. PY2 will also see the startup of a new task responsible for extending the current accounting system to encompass the new resource types that will join the production infrastructure such as cloud and high performance computing resources and to evolve towards an economic and billing system needed for the EGI's self-sustainability.



10 REFERENCES

R 1	MS702: Establishing the Operational Tools Product Team https://documents.egi.eu/document/52
R 2	EGI wiki JRA1 section https://wiki.egi.eu/wiki/WP7-jra1
R 3	MS402: Deploying Software in the EGI production Infrastructure https://documents.egi.eu/document/53
R 4	MS503: Software Provisioning Process https://documents.egi.eu/document/68
R 5	EGI RT System Main Page https://rt.egi.eu/rt/
R 6	OTAG wiki page https://wiki.egi.eu/wiki/OTAG
R 7	USAG wiki page https://wiki.egi.eu/wiki/USAG
R 8	OTAG meetings agenda page https://www.egi.eu/indico/categoryDisplay.py?categId=4
R 9	Symfony open source PHP framework http://www.symfony-project.org/
R 10	MS701: Operations Portal Workplan https://documents.egi.eu/document/39
R 11	Operations Portal wiki page https://wiki.egi.eu/wiki/Operation_Portal
R 12	Operations Portal Release Notes Browser http://operations-portal.egi.eu/aboutportal/releaseNotesBrowser
R 13	Pakiti home page: http://pakiti.sourceforge.net/
R 14	PQ1 GGUS release notes https://gus.fzk.de/pages/releasenotes/release2010-05-19.html https://gus.fzk.de/pages/releasenotes/release2010-06-23.html https://gus.fzk.de/pages/releasenotes/release2010-07-21.html
R 15	PQ2 GGUS release notes https://gus.fzk.de/pages/releasenotes/release2010-09-29.html https://gus.fzk.de/pages/releasenotes/release2010-10-27.html
R 16	GOCD B4 Rollout plan https://www.egi.eu/indico/getFile.py/access?contribId=2&resId=0&materialId=0&confId=63
R 17	GGUS Ticket and Matrix used to track GOCD B4 dependency issues https://gus.fzk.de/ws/ticket_info.php?ticket=5914



	https://www.egi.eu/indico/getFile.py/access?contribId=2&resId=1&materialId=0&confId=63
R 18	GOCDDB4 Feedback from NGIs and users http://goc.grid.sinica.edu.tw/gocwiki/GOCDDB4_feedback
R 19	GOCDDB wiki page https://wiki.egi.eu/wiki/GOCDDB
R 20	OTAG agenda main page https://www.egi.eu/indico/categoryDisplay.py?categId=4
R 21	APEL accounting system wiki page https://wiki.egi.eu/wiki/APEL
R 22	Details on the first three SAM Updates https://tomtools.cern.ch/jira/browse/SAM/fixforversion/10027 https://tomtools.cern.ch/jira/browse/SAM/fixforversion/10044 https://tomtools.cern.ch/jira/browse/SAM/fixforversion/10046
R 23	Procedure to validate NGI SAM Instances https://twiki.cern.ch/twiki/bin/view/EGEE/ValidateROCNagios
R 24	SAM ATP component wiki page https://twiki.cern.ch/twiki/bin/view/LCG/ATP (old page) https://tomtools.cern.ch/confluence/display/SAM/ATP (new confluence page)
R 25	SAM ATP VOFeeds wiki page https://twiki.cern.ch/twiki/bin/view/Main/ATPVFeeds
R 26	SAM Updates release notes pages https://tomtools.cern.ch/confluence/display/SAMDOC/Release+Notes
R 27	SAM Setup for ARC Services https://tomtools.cern.ch/confluence/display/SAM/SAM+setup+for+ARC+services
R 28	The New EGI Metrics Portal Development Roadmap http://www.egee.cesga.es/Metrics_roadmap/DMS0.1_EGI_Metrics_Portal_Roadmap_V1.doc
R 29	AMB face to face meeting in March 2011 agenda page https://www.egi.eu/indico/conferenceDisplay.py?confId=351
R 30	Documentation on Message Broker ACLs configuration https://tomtools.cern.ch/confluence/display/MIG/Message+Broker+ACLs
R 31	Operational tools quick links https://wiki.egi.eu/wiki/Operational_tools_information
R 32	UNICORE Integration Task Forge wiki and agenda pages https://wiki.egi.eu/wiki/UNICORE_integration_task_force https://www.egi.eu/indico/categoryDisplay.py?categId=49
R 33	OPS Monitor at CERN URL https://ops-monitor.cern.ch/nagios/
R 34	ACE Component wiki page https://twiki.cern.ch/twiki/bin/view/LCG/ACE



R 35	GSTAT Tool Main Page http://gstat.egi.eu/
R 36	GridView Tool Main Page http://gridview.cern.ch/GRIDVIEW/same_index.php
R 37	Operational Tools NGIs Deployment Plans https://wiki.egi.eu/wiki/Operations_tools_deployment_plans
R 38	MS703: Operational Tools regionalisation Workplan https://documents.egi.eu/document/107
R 39	EGI Helpdesk main documentation page https://gus.fzk.de/pages/docu.php
R 40	Lavoisier Web Service home page http://grid.in2p3.fr/lavoisier/ http://grid.in2p3.fr/software/lavoisier2/features.html
R 41	GOCDDB Regional Module Documentation Page https://wiki.egi.eu/wiki/GOCDDB/Regional_Module_Technical_Documentation
R 42	GOCDDB Regionalisation Page https://wiki.egi.eu/wiki/GOCDDB/Release4/Regionalisation
R 43	GOCDDB Regional Module Development wiki Page https://wiki.egi.eu/wiki/GOCDDB_Regionalisation_Plans
R 44	JRA1 indico category page https://www.egi.eu/indico/categoryDisplay.py?categId=14
R 45	OMB wiki page https://wiki.egi.eu/wiki/OMB
R 46	Regionalisation Task Force wiki Page https://wiki.egi.eu/wiki/Regionalization-1
R 47	Open Grid Forum main webpage http://www.ogf.org/
R 48	International Symposium on Grid and Clouds 2011 web page http://event.twgrid.org/isgc2011/index.html
R 49	The NAGIOS System Homepage http://www.nagios.org/
R 50	ActiveMQ message system homepage: http://activemq.apache.org/
R51	Operational Tools session at the EGI Technical Forum in September 2010 https://www.egi.eu/indico/sessionDisplay.py?sessionId=20&confId=48#20100916
R52	Security Monitoring Group: https://wiki.egi.eu/wiki/EGI_CSIRT:SMG
R53	The STOMP protocol homepage http://stomp.codehaus.org/



R54	The Open Grid Forum Usage Record Working Group http://www.gridforum.org/gf/group_info/view.php?group=ur-wg
R55	The Bazaar Tool Homepage http://www.plgrid.pl/en/our_offer/tools/user_and_administrator_tools/bazaar
R56	The SOAP Protocol specification http://www.w3.org/TR/2007/REC-soap12-part0-20070427/
R57	The GGUS Report Generator https://gus.fzk.de/stat/stat.php