

# EGI-InSPIRE

## ROADMAP FOR THE MAINTENANCE AND DEVELOPMENT OF THE DEPLOYED OPERATIONAL TOOLS

### EU MILESTONE: MS708

---

Document identifier:	EGI-962-final.doc
Date:	<b>27/04/2012</b>
Activity:	<b>JRA1</b>
Lead Partner:	<b>KIT</b>
Document Status:	<b>FINAL</b>
Dissemination Level:	<b>PUBLIC</b>
Document Link:	<a href="https://documents.egi.eu/document/962">https://documents.egi.eu/document/962</a>

---

#### Abstract

This milestone document records the planned technical changes for the operational tools and the use cases they are designed to support

## I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See [www.egi.eu](http://www.egi.eu) for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See [www.egi.eu](http://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	Torsten Antoni	KIT / JRA1	23.02.2012
Reviewed by	<b>Moderator:</b> Nuno L. Ferreira <b>Reviewers:</b> Feyza Eryol	NA2  NGI_TR	05.03.2012
Approved by	AMB & PMB		27.04.2012

## III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	03.02.2012	First draft; with input from all tool developer teams	Torsten Antoni / KIT
2	17.02.2012	Second Draft; Further Input from tool developers; Added references	Torsten Antoni / KIT
3	18.02.2012	Internal Review	Daniele Cesini / INFN
4	22.02.2012	Further input from tool developers after internal review; Document ready for review	Torsten Antoni / KIT
5	23.03.2012	Changes according to comments from the reviewers	Torsten Antoni / KIT
6	18.05.2012	Minor changes according to second review	Torsten Antoni / KIT

## 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/>.

## 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.

## VII. EXECUTIVE SUMMARY

The development of the deployed operational tools is an on-going activity that concerns the common tools that are currently used to support Grid operations. This work ensures the continuing and correct functionality and interoperation of the tools with the deployed middleware. To monitor this work there is a series of milestones labelled “Roadmap for the maintenance and development of the deployed operation tools”. One is planned for the beginning of each project year. The aim of the milestone is to give an overview of the plans for the developments for the operational tools in the following 18 months, describing the general direction of the development and give estimations of the timeframe for these developments. Of course the plans have to be adapted to changes to the operations model and procedures. Decisions taken by the Operational Tools Advisory Group (OTAG) might also alter the foreseen roadmaps. It is also clear that the level of detail that can be provided concerning the content of specific releases will be larger for releases closer to the time of writing of the milestone than for releases close to the end of the period considered.

The tools in the scope of the milestone are:

- Operations Portal
- GGUS
- GOCDB
- Accounting Repository
- Accounting Portal
- Service Availability Monitoring (including support for messaging)
- Metrics Portal

For each of these tools the current status and their dependencies to other tools are described. This is meant to provide the background on which the new developments and plans will be realised. It can be seen that there are lots of interdependencies between the tools. Some of the tools are of utmost importance for the functioning of the overall operations infrastructure. For those tools sophisticated fail over mechanisms should be in place to make sure that the availability and reliability of the tools can be maximised.

All development teams have detailed plans for the third year and beyond. Nevertheless it is important that the respective advisory bodies (OTAG, USAG) steer and monitor this activity. The requirements coming from the users of the various tools need to be channelled and prioritised by these groups and discussed with the developers, to make sure that changes do not jeopardise the integrity of the interplay between the tools.

## TABLE OF CONTENTS

<b>1</b>	<b>INTRODUCTION .....</b>	<b>6</b>
<b>2</b>	<b>EGI OPERATIONAL TOOLS - STATUS AND DEPENDENCIES .....</b>	<b>7</b>
2.1	Operations Portal.....	7
2.1.1	Current Status .....	7
2.1.2	The regional Instance.....	7
2.1.3	Dependencies.....	9
2.2	GGUS .....	10
2.2.1	Current Status .....	10
2.2.2	Dependencies.....	11
2.3	GOCDB.....	12
2.3.1	Current Status .....	12
2.3.2	Dependencies.....	12
2.4	Accounting Repository .....	13
2.4.1	Current Status .....	13
2.4.2	Dependencies.....	14
2.5	Accounting Portal.....	14
2.5.1	Current Status .....	15
2.5.2	Dependencies.....	15
2.6	Service Availability Monitoring Framework .....	15
2.6.1	Current Status .....	15
2.6.2	Dependencies.....	17
2.7	Metrics Portal .....	17
2.7.1	Current Status .....	17
2.7.2	Dependencies.....	17
<b>3</b>	<b>OPERATIONAL TOOLS ROADMAPS.....</b>	<b>18</b>
3.1	Operations Portal.....	18
3.1.1	Operations Portal Roadmap Summary .....	18
3.2	GGUS .....	18
3.2.1	GGUS Roadmap Summary .....	20
3.3	GOCDB.....	20
3.3.1	GOCDB Roadmap Summary.....	21
3.4	Accounting Repository .....	21
3.4.1	Accounting Repository Roadmap Summary .....	23
3.5	Accounting Portal.....	23
3.5.1	Accounting Portal Roadmap Summary .....	23
3.6	Service Availability Monitoring Framework .....	24
3.6.1	Service Availability Monitoring Roadmap Summary .....	26
3.7	Metrics Portal .....	27
3.7.1	Metrics Portal Roadmap Summary.....	27
<b>4</b>	<b>CONCLUSION.....</b>	<b>28</b>
<b>5</b>	<b>REFERENCES .....</b>	<b>29</b>



## 1 INTRODUCTION

The purpose of this milestone document is to gather and summarise the maintenance and development plans of critical operational tools till the end of project year 3 (PY3 ends April 2013) of the EGI-InSPIRE project. These operational tools are essential to achieve the objective of providing a large scale and resilient pan-European distributed computing infrastructure supporting a diverse range of scientific disciplines.

Section 2 sequentially looks at the role and current status of each of the operational tools. It also reviews the functional dependencies of that tool on other operational tools and on specific data sources.

Section 3 outlines the development plans and release cycles as defined by the corresponding operational tools development teams during PY3.

## 2 EGI OPERATIONAL TOOLS - STATUS AND DEPENDENCIES

### 2.1 Operations Portal

The Operations Portal [R 1] is providing information to various actors (NGI Operations Centres, VO managers, etc.) along with related facilities, such as the VO administration tool, the broadcast and downtime system, the different dashboards (Operations dashboard , security dashboard).

#### 2.1.1 Current Status

The Operations Portal is described in detail in the EGI-InSPIRE milestone document MS705 [R 2].

Since that document was written the development of the central instance has been focused on three modules:

- The UCST VO Management Module :  
This module allows UCST to manage the VO ID card structure, information and to check the VO ID cards registration and update. A new section has been added and summarize all problems detected on the VO ID cards (bad values, bad URLs, missing information, etc.).
- The security dashboard , which permits
  - To collect results from Pakiti and Nagios in order to identify security issues
  - To display on a NGI basis the relevant sites with identified security problems
  - To create tickets against sites
  - To transmit results so as to prevent leakage
  - To compute security metrics based on the numbers gathered
- The VO Module :
  - Add of permalinks for consultation of VO ID Card (by name, or alias, or serial)
  - VOMS declaration and update: the registration / update process of the VOMS server has been reviewed to propose a better quality of information

Eight central releases have been deployed during the 2011 year. The current version is version 2.8.

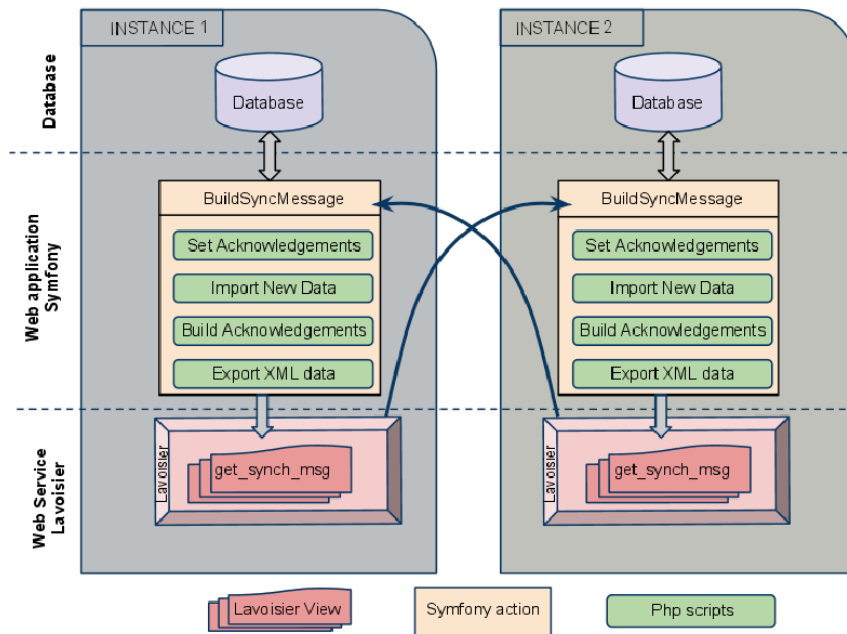
#### 2.1.2 The regional Instance

##### Architecture and synchronization

The application is composed of:

- a web service named Lavoisier [R 3] configured to handle Nagios notifications store and provide data cache from GOCDB and GGUS and to generate metrics reports about the use of the dashboard (alarms raised, tickets handled ...)
- a PHP web application to provide a user interface based on Symfony framework
- a database generated automatically with the web application during the installation

The regional instance is linked with the central instance of Lavoisier; creation, update and deletion of records are synchronized so as not to disrupt global oversight operations. Synchronisation is achieved through REST and SOAP calls and the records are synchronized every 10 minutes by using PHP-scripts. Any problem detected during the synchronisation is reported by email to the webmasters.



**Figure 1: Synchronisation process between Operations Portal instances (central and regional)**

The central and the regional instances (Instances 1 and 2 in Figure 1) have been built on the same model in order to behave in the same way and to be easily interoperable. All regions or NGIs are able to opt for either the central or the regional instance. As shown on the schema (Figure 1) the architecture is exactly the same on a NGI and central instance.

The distributed components are the same, just the configuration changes:

- with PHP code in the application that permits to distinguish a regional instance from a central one
- on the Lavoisier side to filter information and to keep only data relevant to the NGI

## Status

The regional package has been deployed in production in 4 NGIs:

- NGI\_CZ
- NGI\_GRNET
- NGI\_IBERGRID
- NGI\_BY

Three upgrades have been done on the package since the first version of June 2010. The package and the documentation are available through a SVN server [R 4]. The NGI\_UK will set-up a regional instance in the first semester of 2012.

The current work is a refactoring of the package and is focused on:

- The automation of the generation of the package
- The SVN repository organization and the management of the old packages
- The improvement of the synchronization between the central and the regional instances.



The acknowledgements of the synchronization records have been reviewed to avoid blocking the synchronization process, the data have been also encapsulated to avoid problems with special characters.

### 2.1.3 Dependencies

The Operations Portal depends on a number of tools and middleware services to gather status information about the grid (Figure 2). Information gathered from the external sources is cached in Lavoisier (A Data Aggregation and Unification Service) in order to preserve the tool functionality and to increase its robustness in case of problems to retrieve fresh information from the sources.

The Operations Portal depends currently on:

- GGUS
- GOCDB
- SAM PI
- ActiveMQ broker network and Regional Nagios services
- GSTAT
- BDII
- Pakiti
- The Nagios Service for security
- The EGI SSO system
- VOMS servers

The most critical dependency is to GGUS. GGUS tickets are created and updated via SOAP calls. If GGUS is unavailable we can't ensure a correct behaviour of the application. The connection to the EGI RT system is currently under investigation.

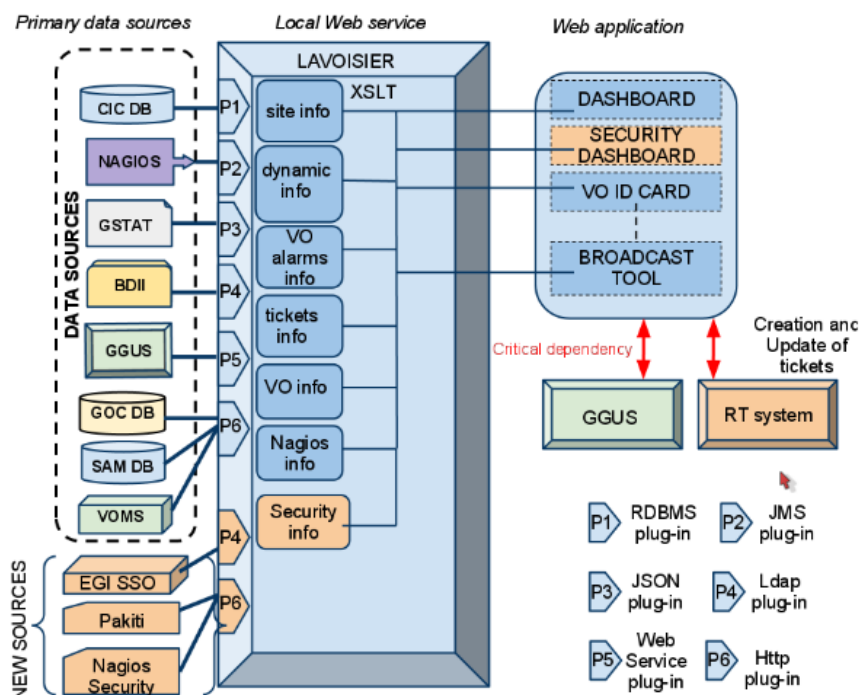


Figure 2: Schematic view of the Operations Portal set-up and its dependencies

## 2.2 GGUS

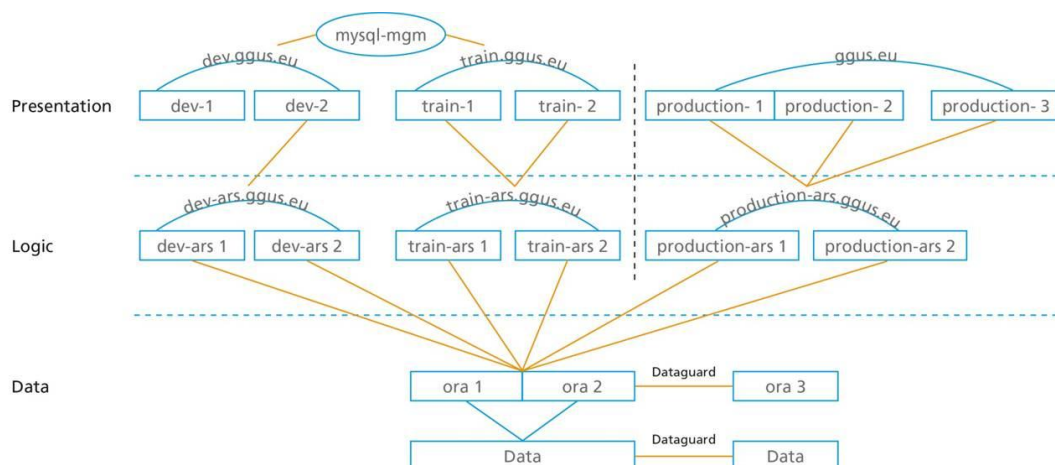
The role of GGUS, as described is that of the EGI helpdesk [R 5, R 6]. Additionally to being the major point of contact for EGI users to report service request and problems, it hosts all project-wide support units and acts as the central integration platform in the distributed support infrastructure consisting of regional and topical helpdesk systems. GGUS is linked into various operational processes through interfaces to other Operational Tools.

### 2.2.1 Current Status

There have been several important topics during the second year. An on-going task is the integration of NGIs into the support infrastructure. This is part of the clearly defined NGI creation process. Thirty-three NGIs support units were created and integrated into GGUS by the end of PY2.

The major development task during PY2 was the refinement of the Technology Helpdesk. This handles all middleware related issues, and acts as a bridge between EGI and its external middleware providers. Currently two workflows are handled in the Technology Helpdesk: (i) the software support workflow for bugs discovered in production; and (ii), the software provisioning workflow. In the latter case, this starts with a release announcement by one of the technology providers, and ends with the acceptance or rejection of the proposed software products. This task is performed under the auspices of the EGI-SA2 activity.

In parallel to the development and adaption work performed in the second year of the project some effort went into further development of the high-availability set-up for GGUS consisting of active-active fail-over components on the data, the logic and the presentation layer (Figure 3). This work will be continued in PY3.




**Figure 3: Schematic view of the three-layered high availability set-up of GGUS**

For further development GGUS depends on the gathering of requirement from its users. These requirements are collected mainly by the OTAG and a newly created GGUS advisory group and prioritised and grouped into releases with the aid of requirements queue of the EGI-RT [R 7].

## 2.2.2 Dependencies

The dependency relationship between GGUS and other tools and other data sources is illustrated in Figure 4. GGUS has direct one-way functional dependencies on:

- GOCDB (site names, email contacts)
- OIM (OSG Information Management System)
- VOMS (rights for Team/Alarmer)



**Figure 4: Schematic view of GGUS dependencies to and from other tools**

Information is regularly retrieved from these systems and stored in a local GGUS cache. In the GOCDB and OIM cases, this information is needed and utilised in facilitating the direct assignment of tickets to sites, e.g. site contact and NGI/ROC affiliation. There are special roles and groups in VOMS defining the membership of the team and alarm ticket submitter groups. Since the retrieved information is stored in a local cache, a failure of these tools ensures that there is a negligible direct impact on GGUS itself. This significantly enhances its availability and reliability. Also, in the case of information retrieval failure, there are no updates to that portion of the cached data. It should also be noted that the data from GOCDB, OIM and VOMS are relatively static, so the use of cached data is pragmatic.

The two-way dependencies can be classified as follows:

- xGUS instances [R 8] provided by and hosted at GGUS,
- Helpdesks interfacing with GGUS.

In the first case, GGUS provides for the following helpdesks:

- Operations Portal
- LHCOPN
- NGI\_AEGIS
- NGI\_CH
- NGI\_DE

If GGUS is in downtime, then these helpdesks are, as a direct consequence, also in downtime. They are all supported and included in the GGUS failover model, which will be finalised at the end of PY3.

In the Second case, GGUS is interfaced with thirteen helpdesk systems hosted by the NGIs and ROCs. During GGUS downtimes they work as stand-alone systems. Tickets are resynchronised afterwards. The interface mechanisms utilises either web services or messaging.

## 2.3 GOCDB

Many aspects of operations rely on the availability of information from NGIs about service nodes, contact details, security contacts, certification status, scheduled downtimes of resources, etc. The GOCDB [R 9] provides all such information through a central database containing sub-databases per operations region, managed by the regions themselves.

### 2.3.1 Current Status

Central GOCDB status [R 10]: GOCDB v4.2 is the current production release. It incorporates new features and some redevelopment of v4.0 modules:

- New MVC logic. This was required to address a number of RT tickets that requested user interface enhancements.
- New database abstractions and a refactored database API. These were required to make the business routines atomic. This has significantly increased reliability and performance.
- Data scoping. This allows non-EGI sites and services to be stored in the central GOCDB.
- New finer grained permissions/role model (scheduled for released in v4.3, Feb 2012).
- Recording of Certification Status History and supporting PI queries.
- Fixing of many long-standing issues.

As a result, operational issues for the central GOCDB have fallen considerably (fewer bugs/GGUS tickets). Patches are progressively rolled out to production. New features are tagged and released under a new version number.

#### Failover status

The central GOCDB database is archived to tape daily. A secure download of the GOCDB database in the form of a dump file is refreshed every two hours for subsequent download and import into the failover instance (hosted at Fraunhofer Institute). The 'goc.egi.eu' DNS switching mechanism has been tested and verified. The failover instance at Fraunhofer needs to be completed.

#### Regionalisation status

Scoping of data has been completed. This allows sites and services to be hosted in the central GOCDB, either in the 'EGI' scope or the non-EGI 'Local' scope. In doing this, NGIs can choose to host their own standalone GOCDB instance, or host their non-EGI sites and services in the central GOCDB. Scoping of EGI and non-EGI data is also a pre-requisite for any future regional-publishing GOCDB. Scoping aside, no other regionalisation work has been undertaken [R 11].

### 2.3.2 Dependencies

As a primary source of information GOCDB doesn't depend on any other tool. For completeness, below is a list of tools that have dependencies on GOCDB:

- Operations Portal and Operations Dashboard (IN2P3)
- ATP (CERN)
- Top-BDII config generator (CERN)
- GridView (CERN)
- NCG (SRCE)
- Regional Nagios (CERN)
- GGUS (KIT)
- GSTAT (ASGC)
- APEL (STFC)
- Accounting Portal (CESGA)

- Metrics Portal (CESGA)
- Grid Resource Bazaar (Cyfronet)

The list is not exhaustive. Some VO specific tools are using the GOCDB programmatic interface (e.g. to feed downtime calendars), and the information may also be used by regional tools (local monitoring, local helpdesks etc.).

## 2.4 Accounting Repository

The EGI accounting repository [R 12] stores information relating to the usage of resources within the EGI production infrastructure. It receives data on individual jobs and summaries of collections of jobs records from information providers, sites and other infrastructures, and exports accounting information to consumers of usage records (for example the accounting portal for visualization (see section 2.5)). This is shown schematically in Figure 5.

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

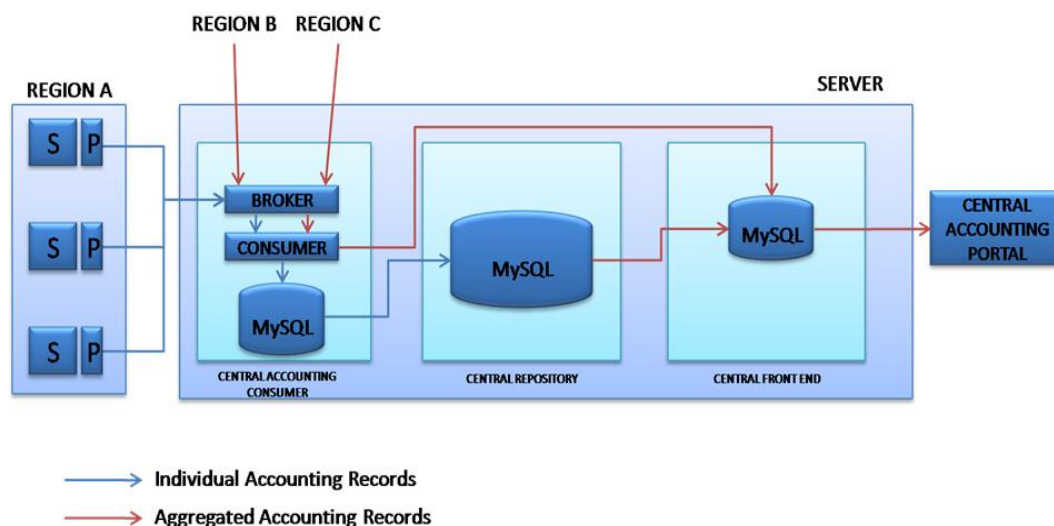
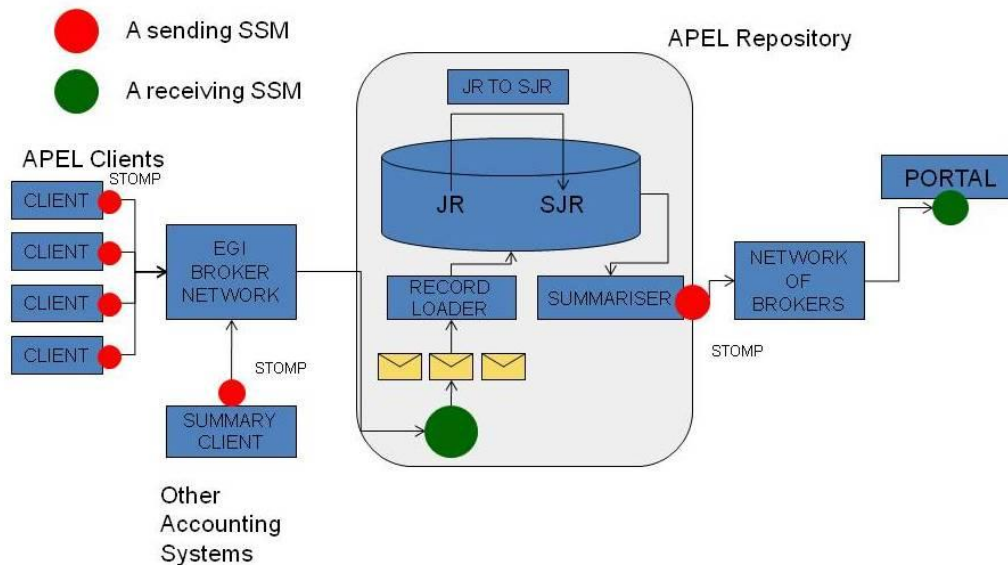


Figure 5: Schematic view of the Accounting Repository

### 2.4.1 Current Status

The APEL Accounting Repository is part way through a major redesign (Figure 6). In part this was to move to the production EGI Messaging Infrastructure and to make the repository suitable for regionalisation. In 2011:

- the database was redesigned to bring it more into line with the OGF UR in addition to improving maintainability and efficiency;
- A new transport layer (SSM) was developed. This uses STOMP and Python and will work on the production EGI Messaging Infrastructure
- SSM and the new repository were tested by the other accounting systems which currently send their data by direct database insertion.



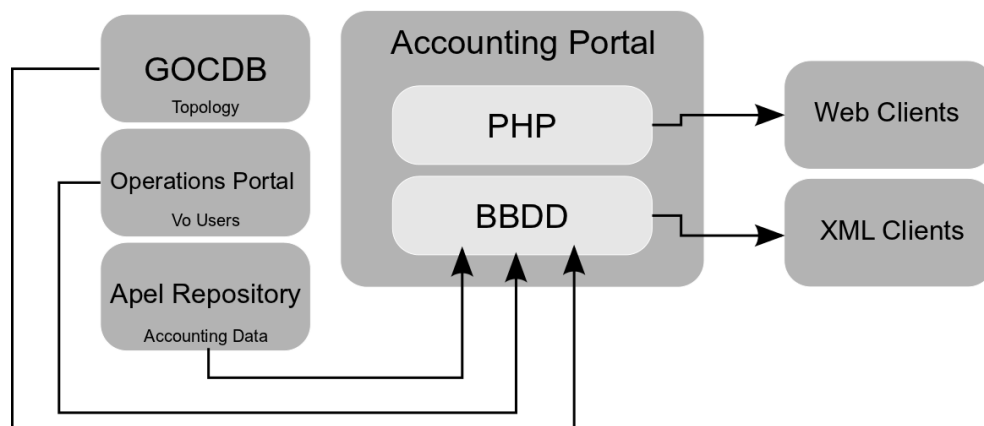
**Figure 6: Outline of the new design of the Accounting Repository**

### 2.4.2 Dependencies

- Messaging Infrastructure: strong dependency but do not foresee problems.
- Accounting Portal: currently it pulls data from APEL by a database query. In 2012 APEL will start to publish summaries by messaging. The portal needs to be ready to receive these messages.

## 2.5 Accounting Portal

The EGI accounting infrastructure is a complex system that involves various sensors in different regions, all publishing data to the central accounting repository. The data is processed, summarized and displayed in the Accounting Portal [R 13], which acts as a common interface to the different accounting record providers and presents a homogeneous view of the data gathered and a user-friendly access to understanding resource utilisation.



**Figure 7: Schematic view of the Accounting Portal dependencies**



### 2.5.1 Current Status

The current production version (v4.0 – Canopus) of the Accounting Portal is available at [R 13]. The regional Accounting Portal is under development (see the roadmap in section 3.5). A beta version is currently available for testing purposes. However, there is no corresponding YAIM plugin available to configure it. A custom installation and configuration script is available.

### 2.5.2 Dependencies

The Accounting Portal has functional dependencies on the following tools (see also Figure 7):

- GOCDB: List of sites and NGIs in production, list of available services in production.
- Operations Portal: VOMS endpoints and VO list.
- Accounting Repository: Accounting records and summarized accounting data.

## 2.6 Service Availability Monitoring Framework

The Service Availability Monitor (SAM) [R 14] is the system that is used to monitor EGI resources within the production infrastructure. It consists of the following components:

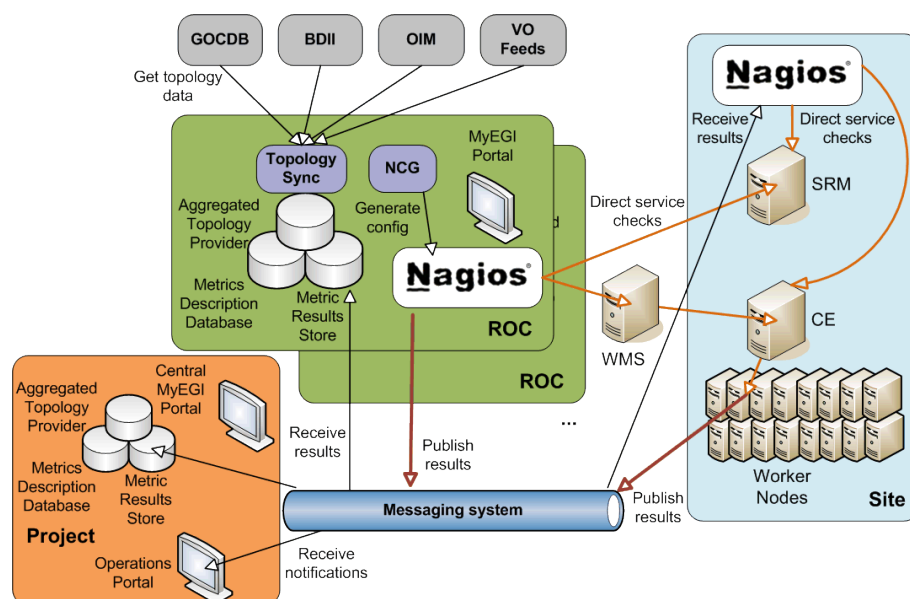
- Probes: a test execution framework (based on the open source monitoring framework Nagios) and the Nagios Configuration Generator (NCG)
- the Aggregated Topology Provider (ATP), the Metrics Description Database (MDDb), and the Metrics Results Database (MRDB)
- a message broker network to publish monitoring results
- a visualization portal (MyEGI)

### 2.6.1 Current Status

The following summarizes the current status per component:

- The test framework: based on the NAGIOS system set up and customized by the NAGIOS Configurator (NCG)
  - The system already fully supports the regionalization and currently all the NGIs are running their respective regional instance.
- Database components: The Aggregated Topology Provider (ATP), the Metric Description Database (MDDb) and the Metrics Result Store (MRS)
  - ATP is currently running in a distributed setting on all the regional instances and synchronizes from multiple source including GOCDB, OSG Information Management System (OIM), GSTAT and VO feeds. In addition, it provides a Web based interface for browsing the synchronized data stored in the database.
  - VO feeds were added as a specific way to introduce and maintain groupings of services as well as support an alternative topological source.
  - MDDb component was deployed to production and provides a way to create profiles binding metrics with service types as well as storing metric descriptions.
  - Metric Result Store (MRS) is currently running in production on all regional instances. It supports computation of service statuses on the regional instances. The central instance aggregates data from NGI instances and provides input to the Availability Computation Engine (ACE).
- Nagios Config Generator (NCG):
  - The NCG is core component of SAM Framework and it is distributed on all regional instances.

- Full integration with the database components (ATP and MRS). All NGI instances are configured to use ATP as topology source.
  - The NCG support probes for ARC, gLite, Globus Toolkit and UNICORE middlewares.
  - Support for failover SAM instances. The failover instance is configured to execute tests but not to raise alarms or report results to the central database.
- Message broker network (Apache ActiveMQ) is used as a reliable and fast message exchange method between operational tools. The production infrastructure uses the PROD message broker network which:
  - consists of 4 geographically separated and under distinct management control brokers
  - is advertised in the grid Information System (BDII) which enables service discovery
- The broker network doesn't fit the regionalization model and is provided as a central service. To facilitate the need for staged rollout and to introduce new clients to the network without risking the operation of the current ones, a totally separated network has been established (TEST network). This network has been successfully used for:
  - The staged rollout of the major update of the messaging service at the end of 2011
  - For minor updates afterwards
  - For the introduction of the APEL needs in the messaging service
- A visualization tool: MyEGI
  - Central monitoring service has been established and put into production. It also contains a production version of MyEGI portal, which is currently providing gridmap, service and metric status views as well as history. The portal is also deployed on the regional instances and replaces the previously used MyEGEE portal.
  - In addition, new implementation of the web service API was provided and is in production (replacing the old SAM-PI). The new API supports XML/JSON serialization and provides backward compatibility with SAM-PI.



**Figure 8: SAM components and their dependencies**



## 2.6.2 Dependencies

Figure 8 provides an overview of the existing SAM components and their dependencies. In particular SAM depends on the following external tools:

- GOCDB
- OIM
- BDII
- GSTAT

The following tools have dependencies on SAM:

- Operations portal (via messaging)
- Metrics portal

## 2.7 Metrics Portal

The Metrics Portal displays a set of metrics that will be used to monitor the performance of the infrastructure and the project, and to track their changes over time. The portal automatically collects all the required data and calculates these metrics before displaying them in the portal. The portal aggregates information from different sources such as GOCDB, GGUS, etc. [R 15].

### 2.7.1 Current Status

The first Production release was open to the public and used for the SA1 metrics of QR6. Now it is being used for QR7, new metrics and HTML and Excel reports for NGIs, quarters and metrics were developed.

### 2.7.2 Dependencies

The Metrics Portal has many dependencies. These include:

- Accounting Portal: To display accounting metrics, most active Vos, Number of submitted jobs, etc.
- BDII: Number of CPUs and Cores in production, online and nearline storage, MPI sites.
- GGUS: Number of tickets created/closed. Tickets response times, Number of tickets created by priority, etc.
- GOCDB: Sites in production, number of countries and NGIs in EGI.
- ACE: Availability and reliability metrics.

## 3 OPERATIONAL TOOLS ROADMAPS

### 3.1 Operations Portal

Two major releases will be deployed in February and March.

The first one will be the finalisation of the development of the security dashboard described in Section 2.1.1. This version will integrate the feedback given by the security group and will be used by Security Operators in the daily security operations.

The second one is the development of a VO oriented dashboard.

This dashboard will provide:

- generic views whose basic entities are sites, tickets and alarms coming from VO SAM tests
- the current number of on-going alarms per site for all sites supporting the VO, and the present number of open tickets per site, for all sites supporting the VO
- the information displayed by an on-going alarm should include details from the failing VO SAM test, and historical data regarding for how long is the alarm active
- the possibility to open tickets against sites
- to compute metrics related to the problems reported on the sites

These 2 releases will be followed by a major upgrade of the regional package.

This upgrade will include the different developments deployed in the 2 previous central releases.

For the long term plan (between May and February 2013) several developments are foreseen.

The priorities will be assessed in agreement with the different stakeholders and the OTAG.

- the refactoring of the regular Operations Dashboard for better ergonomics and efficiency
- a new module used to build availability and reliability charts and reports especially for Top BDII services
- develop a mobile version of the site

#### 3.1.1 Operations Portal Roadmap Summary

Tasks	Planned completion time
Security dashboard : production version	Feb 2012
VO Operations Dashboard : Pilot version	Feb 2012
VO Operations Dashboard : production version	April 2012
Major Upgrade of the regional package	May 2012
Refactoring of the Operations Dashboard	July/Aug 2012
Availability / reliability module	Oct/Nov 2012
Mobile version	March/April 2013

### 3.2 GGUS

#### GGUS Report Generator

The current implementation of the GGUS Report Generator doesn't fit the needs for reporting in EGI. All requirements on reporting have been collected and discussed during a workshop held on

October, 26<sup>th</sup> to 27<sup>th</sup> at KIT in Karlsruhe, Germany. The roadmap for the implementation of the enhanced GGUS Report Generator is as follows:

- Presentation of first prototype on EGI Community Forum
- First version available for the EGI Review mid of the year
- Final version available by end of the year

Implemented features will be:

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

### Integration of further xGUS instances

Several new xGUS instances have to be set up during 2012:

- ROC\_China: work for this xGUS instance has already started but is not finished yet
- NGI\_Slovenia: needs to be set up from scratch
- France Grilles Helpdesk: needs to be set up from scratch
- Maybe some more NGIs

### NGI integration in GGUS

Currently 33 NGIs are integrated in GGUS. The ones that are not yet present are:

- DK      Denmark
- EE      Estonia      Estonian Grid
- LV      Latvia      Latvian Grid
- LT      Lithuania      LitGRID
- NO      Norway      NorGrid
- SW      Sweden      SweGrid

Our current understanding is that those NGIs are all covered by NGI\_NDGF.

Russia still appears as ROC-Russia. This will change once the NGI creation process for Georgia is completed. Then Russia should move from ROC-Russia to NGI\_RU

### Implementation of specific work flows for CSIRT/Security

The CSIRT team is currently evaluating whether they want to use GGUS. If CSIRT will use GGUS, the permissions and access rights schema needs to be adapted to their needs.

### High availability of GGUS

There are two main parts to finalising the GGUS high availability set-up:

- Implementation of auto-switching between Web Front-end Hosts
  - Auto-switching between Web Front-ends
  - Auto-switching between Logic Servers
  - On call duty service integration

- Design disaster recovery plan: processes, policies and procedures related to preparing for recovery or continuation of GGUS after a natural or human-induced disaster.
- Data Base Management System (DBMS) and other high availability questions
  - Intrusion Prevention system configuration for GGUS
  - Improvements into Service Monitoring
  - DBMS on Virtual Environment migration Concept

### Application/Community support

An evaluation of the implementation of web widgets needs to be done. Web widgets allow embedding e.g. the GGUS ticket submit form in dashboards and web pages of user communities. This feature was requested by EGI.eu.

### Alternative Authentication Methods

Currently access to GGUS and xGUS is granted through X.509 certificates and through login and password. Some NGIs have expressed interest in having alternative authentication methods integrated. Shibboleth seems to be one of the most used and most promising federated identity management technologies that could be of interest in the EGI environment. Integrating Shibboleth based authentication in GGUS and xGUS is part of the mid-term roadmap.

### 3.2.1 GGUS Roadmap Summary

Tasks	Planned completion time
GGUS Report Generator (1 <sup>st</sup> prototype)	March 2012
Application/Community support	April 2012 onwards
GGUS Report Generator (1 <sup>st</sup> version)	May 2012
High availability of GGUS (auto-switching)	July 2012
GGUS Report Generator (Production version)	Dec 2012
High availability of GGUS (DBMS)	Dec 2012
Integration of further xGUS instances	On going
NGI Integration	On going
Specific work flows for CSIRT/Security	Still under discussion
Authentication through Shibboleth	To be determined

### 3.3 GOCDB

#### General direction

For the next six months, the general direction of development will focus on; a) continued improvement of the code, including further abstraction of the database and further integration of the new MVC logic; and b) responding to newly emerging requirements for the central GOCDB. To date those include:

- Release of GOCDB v4.3 (Feb), to feature a finer grained permissions and role model (some adjustments to this model are scheduled for agreement by OMB).
- Refine the Virtual Sites proposal with the aim to release Vsites in March/April (v4.4). This will include a new Vsite GOCDB entity, new PI queries and user interface, and potentially new Vsite permissions/roles. This development has already been discussed and agreed in OTAG/JRA1. Some remaining requirements still need to be determined.
- Support for VOs (unconfirmed), such as querying for services according to their supported VO. This requires further discussion within OTAG/JRA1.
- GLUE2 compatibility (unconfirmed), such as new PI queries to render GOCDB data in GLUE2 compatible XML and maybe LDIF Formats. This will most probably require redevelopment of the current XML output module. This requirement needs further discussion within OTAG/JRA1.

### Regionalisation

Since NGIs can now host their local sites and services centrally, a GOCDB regionalisation re-prioritisation exercise is necessary (the outstanding use-case that is addressed by a regional-publishing GOCDB model is to facilitate NGI-specific customisations to non-EGI sites and services).

### 18 Month Roadmap

Post review, if no more high-priority requirements for the central GOCDB emerge, the regionalisation work could commence approximately in the last quarter of 2012. This would form the longer term 18 month roadmap (tentative – this largely depends on the scope of the roadmap items above).

The development roadmap will be discussed and agreed by OTAG and JRA1. Detailed documentation is provided at the GOCDB wiki [R 16].

### 3.3.1 GOCDB Roadmap Summary

Tasks	Planned completion time
Finer grained roles/permissions model	Feb 2012
Virtual Sites	March/April 2012
Redevelopment of XML output module (unconfirmed)	June/July 2012 (tentative)
GLUE2 compatibility (unconfirmed)	July/Aug 2012 (tentative)
Regionalisation reprioritisation exercise	Sept 2012 (tentative)
Support for VOs (unconfirmed)	Depends on requirement/scope of work

### 3.4 Accounting Repository

The vision is of a distributed accounting infrastructure with a central repository that receives several types of usage records and/or summaries from NGIs and other e-infrastructures. Records will originate from: APEL clients; regional instances of APEL, alternative NGI accounting systems; site specific publishers. Several middleware stacks will be supported including ARC, UNICORE and gLite from EMI, and Globus from IGE.

#### JRA1.4

This development task aims to extend the types of usage record collected. This has progressed at different rates.

- **Storage**

EMI have led the development of a storage accounting record (StAR) and we are working with them in the OGF Usage Record Working Group [R 17] to develop this into a standard. Meanwhile it is moving forward as an interim standard in EMI where the storage product teams are preparing to populate records. They will publish using SSM and the APEL repository will receive the records and load them into a database. The release of the EMI products is not scheduled until EMI-3 in May 2013 but this may be negotiable if EGI has a strong requirement. There is an issue of non-EMI storage products (BestMAN, Castor, xrootd) supporting the same standard. The visualisation of storage data in the portal still needs to be planned.

- **Parallel Jobs**

During a revision of the Compute Accounting Record (CAR) which is the original OGF UR, it was confirmed that it was sufficient for parallel jobs to publish the number of compute nodes used by a job and the number of CPUs. It is intended that CPUs contains the number of cores for compatibility with the accepted practise for serial jobs which run on a single core. Unfortunately APEL did not implement these fields, so this is a target for 2012. The repository will be ready to receive them when the database is migrated but will not actually receive the data from most sites until the new APEL client is rolled out. It is not yet understood how best to visualise data from parallel jobs. It is planned to summarise separately for parallel jobs in the repository.

- **Clouds**

The APEL Team are working in the EGI Virtualisation and Clouds Task Force on their accounting scenario. A “straw man” Usage Record starting from the CPU one has been defined and the partners running the various infrastructures are reporting back on how well they can cut them. Testing can start any time but will be more reliable once the new production infrastructure is in place.

By the end of 2012 APEL should:

- be using the production EGI Messaging Infrastructure
- have other accounting systems (SGAS< Gratia, DGAS) publishing by messaging
- doing data transfer to the accounting portal by messaging
- receiving CPU data from Unicore, ARC/JURA, GridSafe(Globus/IGE)
- Receiving additional types of accounting data – storage, clouds
- Using a new APEL client from EMI which publishes using SSM
- Be ready to distribute a regional version of the portal. This will publish using SSM and many of the same summary producer/consumer components developed for the other accounting systems.

### 3.4.1 Accounting Repository Roadmap Summary

Tasks	Planned completion time
New SSM-based infrastructure in production	March 2012
Receiving Test Storage Records	April 2012
Receiving test cloud records	April 2012
Database migration	June 2012
Emi-apel client using SSM released	Nov 2012
Regional distribution of repository	Dec 2012

## 3.5 Accounting Portal

The general direction of the Accounting Portal development is to improve the current code, replace the central Accounting registry connector by a message system and implement the new features requested by the NGIs, VOs, PMB and OMB. New production releases will be available every 6 months but any change will be discussed before within JRA1 and OTAG. The scheduled roadmap for the next 18 months will be available as a DocumentDB space from the Accounting Portal web page.

Version 4.1 of the Accounting Portal will allow seeing the foreign user proportion by VO scope (national, international, etc.). The code base of the regional portal will be improved. The XML endpoints will be generalized and documented.

In version 4.2, the reports will include a view of the contributed CPUs by site, and preliminary support for parallel (MPI) jobs will be added.

JRA1.4 tasks are pending on the development of the new ActiveMQ interface, and slated for the 4.2 version.

### 3.5.1 Accounting Portal Roadmap Summary

Tasks	Planned completion time
Accounting Portal v4.1 <ul style="list-style-type: none"> <li>VO scope views v1.</li> <li>Regional Accounting portal.</li> <li>Generic REST XML endpoints.</li> </ul>	April 2012
Accounting Portal v4.2 <ul style="list-style-type: none"> <li>Contributed CPUs view.</li> <li>Parallel Job accounting.</li> </ul>	Oct 2012



### **3.6 Service Availability Monitoring Framework**

The development will focus on maintenance and bug fixing of the existing components put into production. This would primarily involve MyEGI, Messaging and NCG as core components supported by EGI.

In particular the following developments are planned per component:

#### **MyEGI**

The main effort for the MyEGI portal will be focused on the maintenance of the existing functionality and bug fixing. In particular, the following tasks are planned:

- Maintenance and stability improvements of the central monitoring service.
- Improving the existing documentation and establishing a central point for it.

#### **NCG**

The following tasks are planned:

- Integration of probes for other middleware. Integration of Desktop Grids probes is in progress.
- Implementation of specific SAM instance for monitoring operational tools.
- Redesign of metric configuration. Metrics are currently defined in Perl module and MDDb. In the new version configuration will be moved to config files as part of probe packages and profiles are provided by POEM.
- Migration of grid middleware probes from SAM packages to EMI packages.
- Maintenance and bug fixing.

#### **Messaging**

The main effort at the messaging service activity is to provide secure and reliable and scalable infrastructure. To achieve this goal, four major activities have been identified:

- Security Improvements

The current setup of the broker network service allows both authenticated and unauthenticated connections but the majority of the clients are using unauthenticated connections. This makes it impossible to implement proper ACLs for queues/topics.

In order to authenticate the clients there are 2 options, either use username/password credentials or x509 clientAuth. The possibility of using Grid proxy certificates is still being investigated.

The infrastructure is ready to use authenticated connections but the following steps have to be followed:

- Identify and register the clients
  - Setup credentials for each client
  - Modify client configurations to use authenticated access
- Reliability and availability improvements

In order to consume messages, clients need to specify the endpoint they need to connect to. Currently the majority of the clients choose one of the available brokers and use it without any failover options enabled. STOMP protocol provides failover functionality that can be used by clients either consuming or producing results. In most of the cases, given that the clients are using correctly the STOMP protocol, the only change that would be required on the client side is the connection endpoint.

After the upgrade of broker network software (to ActiveMQ 5.5) at the end of 2011, a new feature is available that allows dynamic failover of the servers. This allows message broker



administrator to dynamically add or remove brokers on the network and request clients to move to another instance. This functionality will be tested on the TEST broker network and if the tests are successful, it will be implemented in the PROD network too.

Another feature that is available currently at the PROD messaging broker network is the Virtual Destinations. This feature minimizes the loss of messages due to disconnections for Topic destinations. Clients can consume from special Queues (Virtual Queues) instead of Topics. Queues keep the messages until the consumer receives them.

- Scalability improvements

The list of the brokers that are connected to the PROD MSG Network is always available through the BDII information system. It is strongly advised that all producers and consumers that use the PROD MSG Network use the BDII in order to find the brokers that are part of the PROD MSG Network at any time. In order not to inquire the BDII every time that a consumer or producer wants to use the PROD MSG Network, information can be cached. The cache must refresh its information at least every day while it is highly recommended to do so in much more frequent periods (i.e. once per hour).

Although Virtual Queues improve reliability, they can have a negative impact on the scalability of the service. Each Virtual Queue that is created for a destination adds a multiplier to the messages that are produced for each message submission, so if one message is sent to a topic with 10 virtual queues, 11 messages will be sent by the broker network. Given that Virtual Queues behave just like the normal Queues, if there is no subscriber to consume the messages, these messages will remain stored at the broker's store forever. In order to eliminate this issue, all messages with an expired TTL will be automatically moved to a special queue (Dead Letter Queue) in order to be cleaned by the system maintenance scripts.

Idle client eviction is another option that has been tested on the TEST Broker Network in order to be enabled at the PROD Broker Network. Currently brokers are keeping connections alive until the clients close them. This can lead to a large number of idle connections (either by producers that don't send anything or by consumers that listen on queues/topics where there is no message). Given that each connection to a broker is consuming resources of the service, idle connections will be evicted. Of course the client can reconnect at any time or even send anti-IDLE messages to avoid unnecessary disconnections.

- Best practices

Given the number of programming languages and libraries that are available for connecting to the broker network it is impossible to cover all cases. The aim will be to cover the following functions:

- Broker discovery
- Message producer
- Message consumer

by:

- Setting the requirements of the each function
- Having the function implemented in at least one programming language as a reference

Write access to the code repository could be enabled for others to contribute.

### 3.6.1 Service Availability Monitoring Roadmap Summary

Tasks	Planned completion time
Messaging: Implementation of credential sharing procedure between brokers	Feb 2012
Messaging: Usage of virtual queues by clients	Feb 2012
Messaging: Creation of the best-practices repository containing at least one of the described functions	Feb 2012
NCG: Redesign of metric configuration and integration with POEM.	March 2012
NCG: Implementation of specific SAM instance for monitoring operational tools.	March 2012
NCG: Integration of Desktop Grids probes.	March 2012
Messaging: First round of client identification	March 2012
Messaging: First round of communication with clients to setup credentials / request of clients to migrate to authenticated connections	March 2012
Messaging: Test of (dynamic) failover functionality	March 2012
MyEGI: Maintenance and stability improvements of the central monitoring service	June 2012
MyEGI: Improving the existing documentation (adding MyEGI user guide).	June 2012
NCG: Migration of grid middleware probes from SAM packages to the EMI packages.	June 2012
Messaging: Third round of client identification, check which instances are still using unauthenticated connections	Sept 2012
Messaging: Third round of communication with clients to setup credentials / request of clients to migrate to authenticated connections	Sept 2012
Messaging: Request to EGI OMB to approve authenticated only connections to PROD message broker network	Sept 2012
Messaging: Enforce authentication to all clients if OMB approves it	Dec 2012
Messaging: Enforce authorization rules if OMB approves it	Dec 2012

### 3.7 Metrics Portal

With the second release of the metrics Portal, the manual metrics will be expanded and refined. New reports with Excel support will be offered for NGI, quarter and metric. The views will be enhanced and optimized, GGUS metrics will be improved and Availability/Reliability metrics added.

The third release will expand and improve the manual metrics. New customized reports with Excel support will be offered. The views will be enhanced and optimized, GGUS metrics will be improved and A/R metrics expanded.

New sets of metrics are included each month from different data providers.

#### 3.7.1 Metrics Portal Roadmap Summary

New Tasks	Planned completion time
Second Metrics Portal release: <ul style="list-style-type: none"><li>• Manual Metrics for the activity leaders v2.</li><li>• QR reports generator v2.</li><li>• Accounting Metrics views v2.</li><li>• GGUS Metrics v2.</li><li>• Grid Information System Metrics v2.</li><li>• Availability/Reliability Metrics v2.</li></ul>	Feb 2012
Third Metrics Portal release: <ul style="list-style-type: none"><li>• Manual Metrics for the activity leaders v3.</li><li>• QR reports generator v3.</li><li>• Accounting Metrics views v3.</li><li>• GGUS Metrics v3.</li><li>• Grid Information System Metrics v3.</li><li>• Availability/Reliability Metrics v3</li></ul>	Aug 2012



## 4 CONCLUSION

This document shows that the deployed operational tools are all interdependent. It is therefore essential that JRA1 ensures that the release process for all these tools is well defined and takes these dependencies into account when testing new releases of the tools. This is done in regular JRA1 and OTAG meetings, where the release procedures and the specific releases for all these tools are presented and discussed, to make sure that this complex system of interlinked tools doesn't break.

All the operational tools have detailed roadmap and development plans for the third year of the project and beyond. The plans are subject to change as the operational procedures are evolving, and the tools have to be adapted to any such changes. To channel the requirements coming in from the users of the tools, it is important that the advisory body, the OTAG, regularly assesses and prioritises the requirements, making use of the requirements queue of the EGI-RT, and check for dependencies with other tools that might be affected by those new developments.

## 5 REFERENCES

R 1	Homepage of the Operations Portal <a href="http://operations-portal.egi.eu/">http://operations-portal.egi.eu/</a>
R 2	MS705 Operations Portal Roadmap: <a href="https://documents.egi.eu/document/525">https://documents.egi.eu/document/525</a>
R 3	Lavoisier Home page : <a href="http://grid.in2p3.fr/lavoisier">http://grid.in2p3.fr/lavoisier</a>
R 4	SVN repository for Operations Portal Regional Package : <a href="https://cvs.in2p3.fr/operations-portal/regional/trunk/">https://cvs.in2p3.fr/operations-portal/regional/trunk/</a>
R 5	Homepage of the EGI Helpdesk <a href="http://helpdesk.egi.eu/">http://helpdesk.egi.eu/</a>
R 6	MS410 EGI Helpdesk and the NGI support units: <a href="https://documents.egi.eu/document/522">https://documents.egi.eu/document/522</a>
R 7	EGI operational tools requirements collection: <a href="https://rt.egi.eu">https://rt.egi.eu</a>
R 8	xGUS homepage: <a href="http://xgus.scc.kit.edu">http://xgus.scc.kit.edu</a>
R 9	Homepage of the GOCDB <a href="https://goc.egi.eu/">https://goc.egi.eu/</a>
R 10	GOCDB development plans: <a href="https://wiki.egi.eu/wiki/GOCDB/Release4/Development">https://wiki.egi.eu/wiki/GOCDB/Release4/Development</a>
R 11	GOCDB regionalisation plans: <a href="https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation">https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation</a>
R 12	MS 706 Operational Tools Accounting Work Plan: <a href="https://documents.egi.eu/document/531">https://documents.egi.eu/document/531</a>
R 13	EGI Accounting Portal: <a href="https://accounting.egi.eu">https://accounting.egi.eu</a>
R 14	Central SAM documentation <a href="https://tomtools.cern.ch/confluence/display/SAMDOC/Home">https://tomtools.cern.ch/confluence/display/SAMDOC/Home</a>
R 15	EGI Metrics Portal: <a href="https://metrics.egi.eu">https://metrics.egi.eu</a>
R 16	GOCDB documentation: <a href="https://wiki.egi.eu/wiki/GOCDB">https://wiki.egi.eu/wiki/GOCDB</a>
R 17	OGF Usage Record Working Group: <a href="http://www.ogf.org/gf/group_info/view.php?group=ur-wg">http://www.ogf.org/gf/group_info/view.php?group=ur-wg</a>