

EGI-InSPIRE

ROADMAP FOR THE MAINTENANCE AND DEVELOPMENT OF THE DEPLOYED OPERATIONAL TOOLS

EU MILESTONE: MS707

Document identifier:	EGI-doc-523-V2.doc
Date:	09/08/2011
Activity:	SA 1
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/523

Abstract

The development of the deployed operational tools is an ongoing activity that concerns the common tools that are currently used to support Grid operations. This work ensures the continuing and correct functionality and interoperability of the tools with the deployed middleware.

This milestone document is a public report giving a short overview of the current status of the operations tools and describing the roadmaps for all the deployed operational tools over the next 18 months, starting from PM 15, defining release and deployment dates.

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	Torsten Antoni	KIT / SA 1	17/05/2011
Reviewed by	Moderator: Enol Fernandez Reviewers: John Walsh	CSIC/SA2 TCD/SA1,SA3	27/07/2011
Approved by	AMB & PMB		09/08/2011

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	29/06/2011	First draft – Compiled from input from all developers teams of the central tools	Torsten Antoni / KIT
2	29/07/2011	Second draft including comments and corrections by the reviewers	Torsten Antoni / KIT
	03/08/2011	Minor corrections from second round of external 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/>.

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 development of the deployed operational tools is an ongoing 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. 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
- 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 second 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 don't jeopardise the integrity of the interplay between the tools.

TABLE OF CONTENTS

1	INTRODUCTION.....	6
2	EGI OPERATIONS TOOLS - STATUS AND DEPENDENCIES.....	7
2.1	Operations Portal.....	7
2.1.1	Current Status.....	7
2.1.2	Dependencies.....	7
2.2	GGUS.....	8
2.2.1	Current Status.....	8
2.2.2	Dependencies.....	9
2.3	GOCDB.....	10
2.3.1	Current Status.....	10
2.3.2	Dependencies.....	10
2.4	Accounting Repository.....	11
2.5	Accounting Portal.....	11
2.5.1	Current Status.....	11
2.5.2	Dependencies.....	11
2.6	Service Availability Monitoring Framework.....	11
2.6.1	Current Status and Regionalization.....	12
2.6.2	Dependencies.....	13
2.7	Metrics Portal.....	13
2.7.1	Current Status.....	13
2.7.2	Dependencies.....	14
3	OPERATIONS TOOLS ROADMAPS.....	15
3.1	Operations Portal.....	15
3.1.1	Tasks inherited from the first year.....	15
3.1.2	New tasks.....	15
3.2	GGUS.....	16
3.3	GOCDB.....	17
3.4	Accounting Repository.....	18
3.4.1	Server Redesign and Development.....	18
3.4.2	Other Developments.....	20
3.5	Accounting Portal.....	20
3.6	Service Availability Monitoring Framework.....	21
3.7	Metrics Portal.....	22
4	CONCLUSION.....	23
5	REFERENCES.....	24



1 INTRODUCTION

The purpose of this milestone document is to gather and summarise the maintenance and development plans of critical operational tools developed during project year 2 (PY2) 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 status of each of the current 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 PY2.

2 EGI OPERATIONS TOOLS - STATUS AND DEPENDENCIES

This section gives a short overview of the current status of all the operations tools in scope of this milestone document and their interdependencies. It describes the foundation of the plans for the further development of the operations tools described in section 3.

2.1 Operations Portal

The Operations Portal is providing information to various actors (NGI Operations Centres, VO managers, etc.) along with related facilities, such as the VO registration tool, the broadcast and downtime system, the periodic operations report submission system, the regional dashboard, etc.

2.1.1 Current Status

The Operations Portal is described in detail in the EGI-InSPIRE milestone document MS705 [R 1].

2.1.2 Dependencies

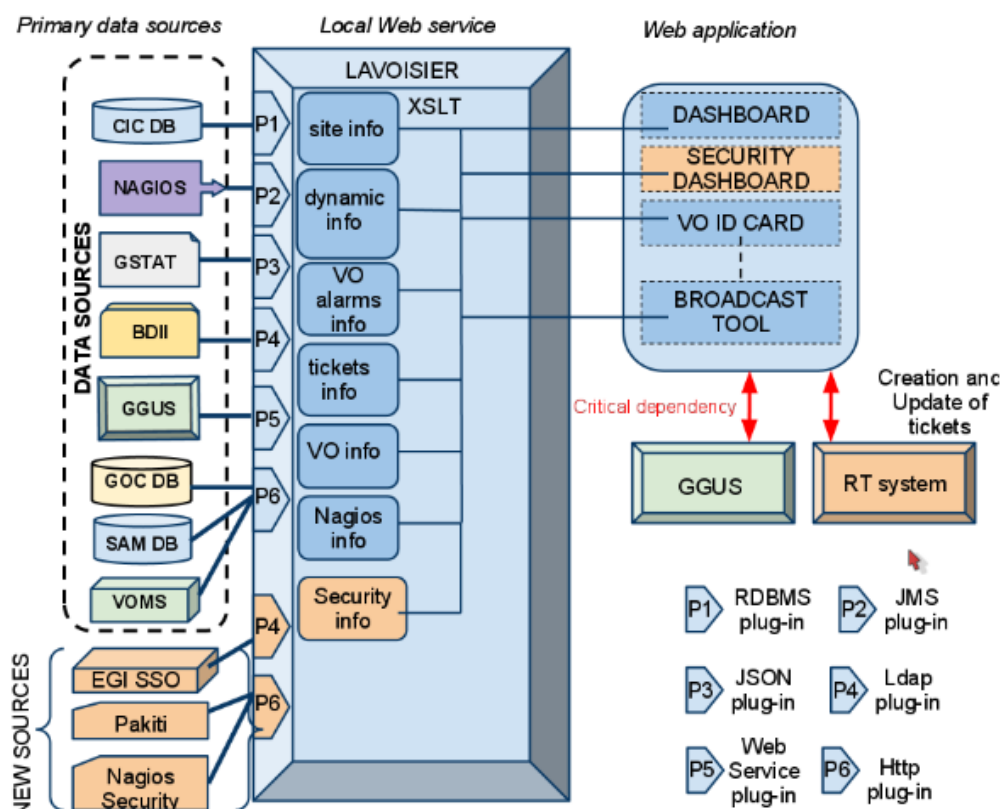


Figure 1: Schematic view of the Operations Portal set-up and its dependencies

The Operations Portal depends on a number of tools and middleware services to gather status information about the grid (Figure 1). Information gathered from the external sources is cached in Lavoisier 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:

- GOCDB
- SAM PI
- ActiveMQ broker network and Regional Nagios services
- GSTAT
- BDII

New dependencies are foreseen in the coming months:

- Pakiti
- The Nagios Service for security
- The EGI SSO system

Services on which the Operations Portal depends in a way that their malfunctioning strongly influences the Operations Portal are GGUS and VOMS.

GGUS tickets are created and updated via SOAP calls. If GGUS is unavailable we can't ensure a correct behaviour of the application. This solution will be extended in the same way with the RT system to create tickets for the security team.

Information is retrieved from the different VOMS servers to feed the Operations Portal User Database. If a VOMS server is unavailable for some time the resulting list of users is not updated but the functionality of the tool is not compromised.

2.2 GGUS

The role of GGUS, as described in MS410 [R 2], is that of the EGI helpdesk. 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 operations tools.

2.2.1 Current Status

The role of GGUS, as described in MS410 [R 2], is that of the EGI helpdesk. Various workflows had to be adapted and details to be changed. There have been several important topics during the first 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-one NGIs support units were created and integrated into GGUS by the end of PY1. However, this represents only approximately 75% of the total number of NGIs, which means that roughly 25% has not yet finished the NGI creation process.

The major development task during PY1 was the creation 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 first year of the project significant effort went into the high-availability set-up for GGUS consisting of active-active fail-over components on the data, the logic and the presentation layer (Figure 2).

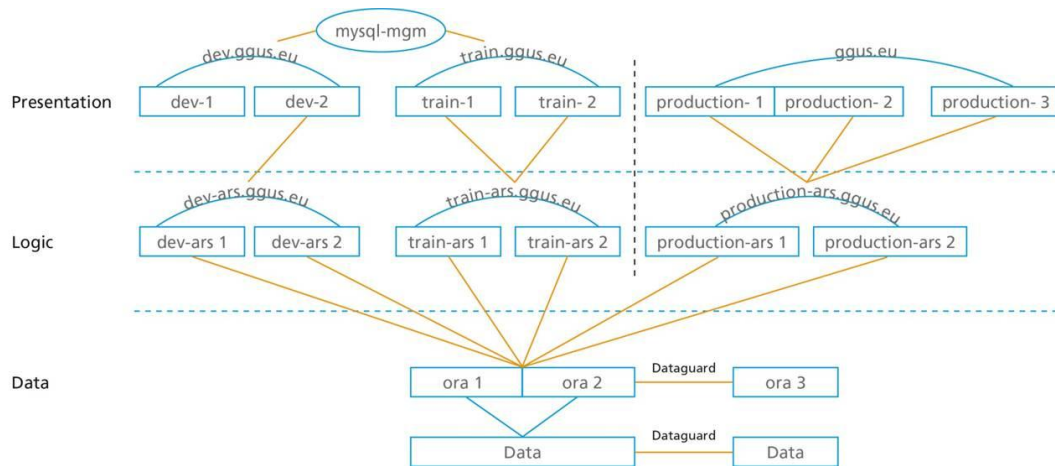


Figure 2: Schematic view of the three-layered high availability set-up of GGUS

For further development GGUS depends on the gathering of requirement from its users. These requirements are collected mainly by the OTAG and the USAG and prioritised and grouped into releases with the aid of requirements queue of the EGI-RT [R 3].

2.2.2 Dependencies

The dependency relationship between GGUS and other tools and other data sources is illustrated in Figure 3. GGUS has direct one-way functional dependencies on:

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

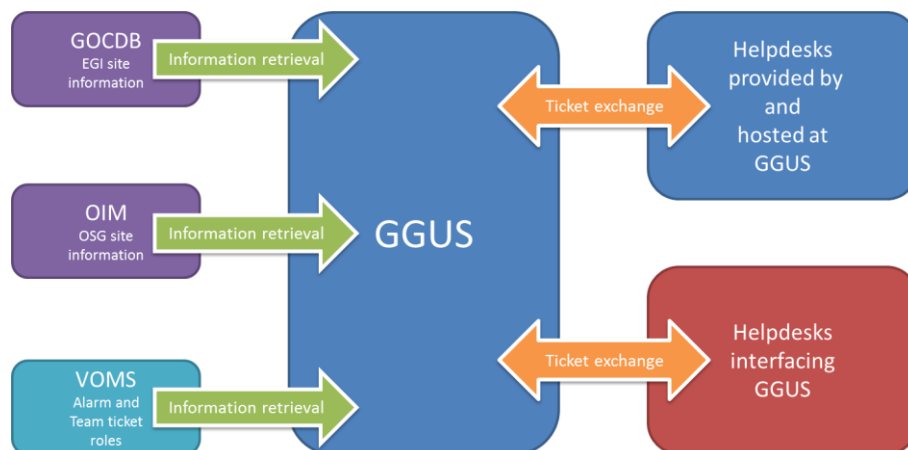


Figure 3: Schematic view of GGUS dependencies to and from other tools

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:

- Helpdesks provided by and hosted at GGUS,
- Helpdesks interfacing with GGUS.

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

- CIC
- 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 the year.

In the Second case, GGUS interfaces 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. Some older ROC helpdesk interfaces are still managed partially by e-mail.

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 provides all such information through a central database containing sub-databases per operations region, managed by the regions themselves.

2.3.1 Current Status

GOCDB4 was released in the last quarter of 2010, and represented a major functional re-write. Following its release a number of new feature requests have been recorded (43 RT tickets) and some design restrictions have emerged. Current work is thus focussed on operational support, bug fixing and re-development to accommodate new (post V4 release) requirements. The service is hosted centrally for EGI by the UK STFC. Patches and new service functions are progressively rolled out to production. A Regional-Standalone version is available for download and customisation (this version does not synchronize with the central GOCDB). A Regional-Synchronizing version that will synchronise its EGI scoped regional data with the central instance is planned but is a longer term development.

Up-to-date information on GOCDB can be found at [R 4] and on the status of regionalisation at [R 5].

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)
- SAMAP (Poznan)
- APEL (STFC)
- Accounting Portal (CESGA)

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 stores information relating to the usage of resources within the EGI production infrastructure. It receives records from information providers, and exports accounting information to consumers of usage records (for example the accounting portal for visualization).

The current status and dependencies of the accounting repository is described in MS706 [R 6].

2.5 Accounting Portal

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

2.5.1 Current Status

The current production version (v3 - Betelgeuse) of the Accounting Portal is available at [R 7]. 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:

- GOCDB: List of sites and NGIs in production, list of available services in production.
- CIC 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) 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 and Regionalization

The following summarises the current status of each SAM component:

- The test framework: based on the NAGIOS system set up and customized by the NAGIOS Configurator (NCG)
 - Database components: The Aggregated Topology Provider (ATP), the Metric Description Database (MDDDB) 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, RSV, 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.
 - MDDDB 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.
 - NCG was integrated fully with all database components (ATP, MRS). All NGI instances are configured to use ATP as topology source.
 - NCG was extended with probes for ARC and Globus Toolkit middlewares.
- Message broker network (ActiveMQ) to publish the monitoring results (PROD MSG broker network)
 - The PROD message broker network consists of 4 brokers. Each broker is geographically separated and under distinct management control
 - Each broker publishes their status into the grid Information system (BDII). This enables service discovery.
 - Due to architectural design, it is not recommended or required to provide regionalised message brokers.
- 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, a 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.

2.6.2 Dependencies

Figure 4 illustrates the existing SAM model, its components and their dependencies. In particular SAM depends on:

- GOCDB
- RSV, OIM
- BDII
- GSTAT

The following components have dependencies on SAM:

- Operations portal (via messaging)
- Metrics portal

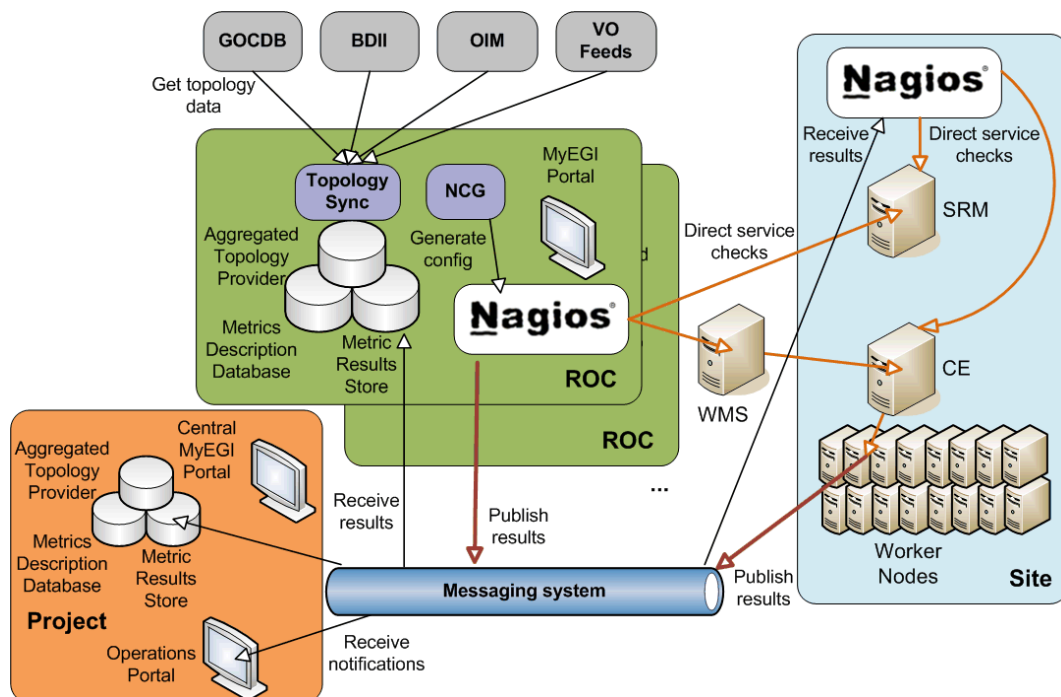


Figure 4: SAM components and their dependencies

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.

2.7.1 Current Status

A new development version of the Metrics Portal [R 8] was developed in the last months of PY1. Manual metrics views are now available in the portal and can be filled directly by the activity leaders. The Metrics Portal is still in development and new metrics will be included in next months before the production release (August 2011).



2.7.2 Dependencies

The Metrics Portal has many dependencies. These include:

- Accounting Portal: To display accounting metrics, most active VO's, Number of submitted jobs, etc.
- BDII: Number of CPUs and Cores in production.
- GGUS: Number of tickets created/closed. Tickets response times, Number of tickets created by priority, etc.
- GOCDDB: Sites in production, number of countries and NGIs in EGI.
- ACE: Availability and reliability metrics.

3 OPERATIONS TOOLS ROADMAPS

This section of the milestone documents collects the roadmaps of all the operations tools in scope, describing the planned main development strands and giving timelines for their implementation where this is possible. Of course the priority of the implementation of the various features is subject to change. The progress of the development (and possible changes in the prioritization) of the operations tools is tracked by advisory bodies, mainly OTAG, but for some tools also by other bodies like USAG in the case of GGUS.

3.1 Operations Portal

The roadmap for the Operations Portal is described in detail in EGI-InSPIRE milestone MS705 [R 1]. We will therefore limit the description here to a short summary. The main tasks for the coming year are the following:

3.1.1 Tasks inherited from the first year

Features	Timeline
Downtime Notification System	May 2011
User Tracking	June 2011
Decommissioning of historical Portal instance	June 2011
Release of Lavoisier 2.0	July 2011
Integration of Lavoisier 2.0	July –December 2011
Enhancements of Lavoisier Programmatic Interface	July-December 2011
Harmonisation of GOCDB and Operations Portal	January 2012 => May 2012

3.1.2 New tasks

Features	Timeline
Security Dashboard : prototype	June 2011
Replacement of the VO specific tests	June 2011
Improvements to filters used by the dashboard	July 2011
Integration of x-GUS and RT system	August 2011
Section with a summary of VOs not respecting the rules	August 2011
Security Dashboard : official release	October 2011
Yearly renewal of VO ID cards	November 2011

3.2 GGUS

The following items are on the roadmap for development in the next 18 months. Of course other tasks will come up in the regular meetings of advisory bodies like the OTAG and the USAG, where requirements are collected from the users of the helpdesk infrastructure.

- Improvements of the Technology Helpdesk
 - The Technology Helpdesk acts as a bridge between EGI and its external technology providers. To make sure that the support staff from the various providers have access to the appropriate tickets, a sophisticated access control model is necessary. This will be implemented in the July 2011 release of GGUS.
 - EGI-SA2 needs to be able to create tickets assigned the “release management” support unit from EMI. This means a cross-link between the software provisioning and bug reporting workflows. These tickets will not be created in the technology helpdesk, but in the EGI-RT therefore some adjustments need to be made on both ends of the GGUS-RT interface. This feature will be part of the September 2011 release of GGUS.

- Redesign of the GGUS Report Generator

The current GGUS Report Generator was developed during EGEE-III to cover specific reporting needs expressed by various EGEE bodies at that time. The first year of EGI-InSPIRE has shown that there is a strong interest in the possibility to create statistics and metrics out of the ticket data. Not all the requested reports can be created with the current report generator.

Currently the requirements from various EGI bodies are being consolidated into a complete list that groups the requirements by complexity and effort needed to implement them. These range from simple additions to the current report generator to features that would require a complete re-design of the tool.

Several options and possible new tools are evaluated at the moment. A new report generator could reach production status by the end of 2011.

- Completing the NGI integration

As mentioned in paragraph 2.2.1 on the current status of GGUS the integration of the NGIs in to the support infrastructure is not yet finished around 10 NGIs are still absent from GGUS, which means, as this is part of the overall NGI creation process, that these NGIs are not yet fully operational. This is an issue that needs to be addressed.

Of the NGIs that are already connected to the EGI helpdesk the majority uses the helpdesk directly without an interface to their NGI helpdesk. As the NGIs mature and the amount of tickets they receive increases a manual interface between the central EGI helpdesk and the NGI helpdesk can constitute a bottleneck in the support process. It is therefore important to help the NGIs to set up helpdesk systems and to interface them to GGUS.

At the EGI Technical Forum 2011 in September there will be a workshop session dedicated to this issue. The goal is to have all NGIs present in the helpdesk before that session and to gradually increase the number of automatic interfaces after the Technical Forum. The next check point could be the EGI Community Forum in spring 2012.

- Integration of GGUS into community & application support processes

For the users the EGI helpdesk should be known as the one place they can submit all their issues to. The EGI support infrastructure then has to make sure that all issues can be and are routed to the correct place.

- To lower the threshold of experience with the EGI infrastructure needed to get in touch with the support teams it is planned to have contact forms on the EGI web pages that when filled in create a ticket in GGUS. General forms will trigger tickets assigned to the UCST. This feature, that will also require some work on the EGI web pages, will be ready for the September 2011 release of GGUS. In a second step service specific contact forms will gradually be implemented that will trigger tickets automatically assigned to the support teams responsible for the respective services.
- To ensure that issues can smoothly be transferred from the “incident management process” to the “problem management process” (as defined in the ITIL framework) an interface between GGUS (incident management) and the EGI-RT (problem management and requirements collection) needs to be implemented. This Interface needs some work also on the RT side. It will be ready for the October release of GGUS.

3.3 GOCDB

Developments for the next 6-10 months are focused on re-factoring and extending the service to cater for newly emerging requirements that cannot be easily implemented with the current system (some nontrivial developments). In order of priority:

- Adding new Service Endpoint Location associations (URLs, WSDLs etc.) for new service types (e.g. Unicore6 services). This requires refactoring of the existing DB cardinality logic and of the XML Output module. This is necessary because the existing XML Output module will only generate flat XML documents that directly map to individual DB entities (i.e. joins/associations between Sites, SEs and Endpoint Locations cannot be represented as hierarchical/nested XML documents). Planned release: June 2011.
- The GUI logic needs to be re-developed. Currently GOCDB4 uses a generic module for drawing a forms and GUI components that uses a single code path for all form based operations. In doing this, the code has become rigid making it difficult to deal with GUI requirements on a per-page basis. This module will be replaced with a more flexible controller per page architecture. This is necessary to address GUI enhancements (approximately half of the GOCDB RT tickets request GUI enhancements). Planned release: August/Sept 2011.
- The role model will need to be extended for finer grained permissions. This is currently emerging as a new EGI requirement. Planned Release: will be determined depending on requirements.
- New tagging logic is required so that new data-scoping rules can be applied to GOCDB entities. In doing this, Sites, Services and other GOCDB data can be identified as either ‘EGI,’ ‘Local’ or ‘some other’ scoped data. This is required for the Regional-Synchronizing GOCDB so that only appropriately scoped ‘EGI’ data is published to the central instance while locally scoped data and other customisations can be excluded. Entity tagging is also relevant for the ‘Virtual Sites’ requirement (i.e. using tags to collate entities into virtual groupings). Planned Release: end 2011.

Subsequent developments for the 6-10 months following the period described in detail above can then focus more upon regionalisation. The development roadmap will be discussed and agreed by OTAG and JRA1. All documentation is provided at the GOCDB wiki [R 9].

3.4 Accounting Repository

There are two strands to the Accounting repository roadmap. The first strand, covered in subsection 3.4.1, is a major redevelopment of the accounting repository service to enable its deployment as a distributable regional accounting repository. The second strand, covered in subsection 3.4.2, includes the development of new functionalities.

3.4.1 Server Redesign and Development

M1.1 DB Schema and APEL Message formats:

Design the database schema for storing all accounting data. Design the message format for publishing different types of records to the APEL accounting repository (work with EMI) [R 10]:

- Job Records (JR): Individual CPU accounting records.
- Summary Job Records (SJR): Monthly summary CPU accounting records per site, VO and user.
- Sync Records (SR): Summary record with number of jobs per site for calculating Nagios-APEL-sync tests.

This work is done but is included here for completeness.

M2. SJR Consumer in Pre-Production:

A consumer for Summary Job Records is installed centrally to be used by the non-APEL repositories (SGAS, DGAS, OSG, etc.) to develop/test their summary publishing tools. Records published through this interface will not be transferred to the production database at this point.

A publisher (Stomp Secure Messenger) is provided for use or as an example for any other region to implement their own publishing software.

A testing infrastructure is installed (in parallel) for development/testing by external clients.

6th June

M3. SJR Consumer and new APEL Central database in Production:

The consumer for Summary Job Records is ready for the non-APEL clients to publish their production data (only summary data). These records will be merged with the data from the APEL production system for display in the Accounting Portal.

This step is blocked by M6. The new database cannot be put into production and the job record data migrated until the non-APEL publishers are ready to publish by the new method.

This should be ready very soon after M2 but enough time will be given for all non-APEL clients to migrate to this new interface before closing the old direct insertion interface when the old database is migrated.

M4. JR Consumer deployed:

The APEL server is able to receive Job Records (JR) from the messaging system and insert them into the APEL new database.

End of June – allows testing of new EMI-APEL client and any other JR publishers.

M5. Backend record processing completed:

So far the new server receives JR and SJR and stores them. This milestone involves all the post-processing to create SJR from JR and publish them onwards. This is required new functionality that will be used by the regionalized package. More immediately it is how data will be published to the accounting portal.

At this stage the new server is complete and ready to receive JR and SJR from the production AMQ via new publishers.

End of August

M6. New Database accepting updates from legacy AMQ:

Records received through the old APEL broker/SSL interface are redirected as new messages to the network of brokers so they can be received by the new STOMP system. All the accounting data is now stored in the new database and its only source of incoming data is the STOMP interface through the production network of brokers.

The data from the old APEL database are migrated into the new database.

Dependencies:

- The interface between the APEL server and the Accounting Portal has been moved to ActiveMQ messaging, using the same software and interfaces as at the client level taking data produced in M5.
- All the old direct database insertion publishers have migrated to AMQ.

M7. APEL Publisher modified and released:

A new release of EMI-APEL will enable end sites to start publishing directly into the new service.

Mid-September

This is an EMI milestone shown here for completeness.

M8. Deployment of new EMI-APEL:

This is the corresponding EGI SA1 milestone to M7. This will start right away but will take many months to complete.

M9. Regional Accounting Server packaged and released:

The Regional Accounting Server is released for the NGIs to install, together with a migration plan for them to follow. This could start sooner but preferable to have reasonable experience of the central service in production first.

End of December

M10. Deployment and migration completed:

All sites publishing JR directly to the central repository have migrated to new EMI-APEL and old consumer and private broker network can be closed.

3.4.2 Other Developments

Various other functionalities are intended to be supported, but no timescales have yet been set. These will be developed over the next few months after gathering requirements and input from various stakeholders.

Other accounting record types:

After the completion of task M4, an alternative Usage Record consumer could be developed that receives accounting for other resources (e.g. storage, clouds, applications, etc.) and stores them in a database. This would be sufficient for testing new accounting producers. Production use would require M5 and additional work from the accounting portal to process and display them.

RUS Interface:

OGF-RUS (or more recently RUPI) [R 11] is a standard web service interface that receives job usage records. It is desirable to offer a standard-compliant interface to the central repository. Standardisation reduces the work required to support alternative accounting systems used by the NGIs.

Alternative Database Interface

Although the general requirement is to pass on only summary information we have had requests from projects etc. who want/need copies of their own data at the job level. Provided all the legal issues could be solved it would be expedient to offer a general authorized database query interface which would allow for more complicated extractions. For example an individual may extract records only for a particular VO.

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 and NA3. 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. Scheduled dates:

- 10/31/11 Accounting Portal v4.0 (Canopus):
 - New charts based on pChart framework.
 - ActiveMQ connector.
- 04/30/12 Accounting Portal v4.1 (Deneb):
 - VO scope views v1.
 - Regional Accounting portal.
 - Custom XML endpoint for NA3.
- 10/31/12 Accounting Portal v4.2 (Electra):
 - Contributed CPUs view.
 - Parallel Job accounting.

3.6 Service Availability Monitoring Framework

A general direction of 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 component supported by EGI.

In particular the following developments are planned per component:

- MyEGI
 - Improving the existing functionality by adding more features to gridmap view and trying to establish the same level of functionality as provided by the previous gridmap service. This would also involve addressing the issues with FQANs and profiles.
 - Improving stability of the Web Service API.
 - Maintenance and stability improvement of the central monitoring service.
 - Improving the existing documentation and establishing a central point for it.
- NCG
 - Integration of probes for other middleware. Integration of UNICORE probes is in progress. New probes will be integrated when provided by middleware developers.
 - Support for failover SAM instance. The failover instance will be configured to execute tests but not raise alarms or report results to the central database.
 - Redesign of metric configuration. Metrics (tests) are currently defined in Perl module and MDDB. In the new version configuration will be moved to configuration files as part of probe packages.
- Messaging
 - Messaging in EGI-InSPIRE covers only configuration developments thus the roadmap is communicated with CERN as the EMI representative for MSG release
 - Security
 - Secure all broker interfaces
 - Incrementally restrict access to queues and topics only to the expected consumers
 - Accounting
 - Currently there is no accounting information for the broker network
 - Investigate solutions to provide accounting information regarding the usage of the brokers
 - Provide a testing/development broker network for operation tools to test their applications before requesting access to the production network

New releases are delivered in an incremental fashion, with a release every 3-4 weeks. These will be a mix of bug fixes and new feature development. The prioritisation of the above mentioned development plan is currently on-going. It is planned to publish them amongst other important information related to SAM at a central point of documentation, that is currently being established and a working version of which is running at [R 12].

3.7 Metrics Portal

The objective of the Metrics Portal development is to release a production version in August 2011. New sets of metrics are included each month from different data providers. The scheduled dates for the Metrics Portal are:

- 31/08/11 First Metrics Portal production release. It will include:
 - Framework based on Django.
 - Manual Metrics for the activity leaders v1.
 - QR reports generator v1.
 - Authentication access based on user certificate or SSO.
 - Accounting Metrics views v1.
 - GGUS Metrics v1.
 - Grid Information System Metrics v1.
 - Availability/Reliability Metrics v1.
- 28/02/12 Second Metrics Portal release:
 - Manual Metrics for the activity leaders v2.
 - QR reports generator v2.
 - Accounting Metrics views v2.
 - GGUS Metrics v2.
 - Grid Information System Metrics v2.
 - Availability/Reliability Metrics v2.
- 31/08/12 Third Metrics Portal release:
 - Manual Metrics for the activity leaders v3.
 - QR reports generator v3.
 - Accounting Metrics views v3.
 - GGUS Metrics v3.
 - Grid Information System Metrics v3.
 - Availability/Reliability Metrics v3.



4 CONCLUSION

This document shows that the deployed operational tools are all interdependent. It is therefore essential that JRA1 ensures (through the OTAG) that the release process for all these tools is well defined and takes these dependencies into account when testing new releases of the tools. The release procedure for all these tools should at least be centrally monitored, 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 second 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 bodies, USAG and OTAG, regularly assesses and prioritises the requirements, making use of the requirements queue of the EGI-RT, and checks for dependencies with other tools that might be affected by those new developments.

5 REFERENCES

R 1	MS705 Operations Portal Roadmap https://documents.egi.eu/document/525
R 2	MS401 EGI Helpdesk and the NGI support units https://documents.egi.eu/document/522
R 3	EGI operational tools requirements collection https://rt.egi.eu
R 4	GOCDDB development plans https://wiki.egi.eu/wiki/GOCDDB/Release4/Development
R 5	GOCDDB regionalisation plans https://wiki.egi.eu/wiki/GOCDDB/Release4/Regionalisation
R 6	MS706 Accounting Roadmap https://documents.egi.eu/document/531
R 7	EGI Accounting Portal http://accounting.egi.eu
R 8	EGI Metrics Portal (Development Version) https://metrics-devel.egi.cesga.es/
R 9	GOCDDB documentation https://wiki.egi.eu/wiki/GOCDDB
R 10	Documentation of the APEL message format https://wiki.egi.eu/wiki/APEL/MessageFormat
R 11	OGF-RUS working group http://forge.gridforum.org/projects/rus-wg
R 12	Central SAM documentation (working version) https://tomtools.cern.ch/confluence/display/SAMDOC/Home