EGI-InSPIRE

OPERATIONAL TOOLS REGIONALIZATION STATUS

EU MILESTONE: MS401

Document identifier:	EGI-MS401-final
Date:	**05/11/2010**
Activity:	**SA1**
Lead Partner:	**INFN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/48

Abstract

This document assesses the current status of the regionalisation of the operation tools and prioritises the remaining work that needs to be undertaken by WP7.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	L. Gaido	INFN	20/07/2010
Reviewed by	**Moderator:** Steve Brewer **Reviewers:** G. Misurelli, M. David	EGI.eu INFN	28/06/2010
Approved by	**AMB & PMB**		22/07/2010

III. DOCUMENT LOG

ISSUE	Date	Comment	Author/Partner
0.1	15.6.2010	Document structure and TOC	L. Gaido/INFN
0.2	21.6.2010	First partial draft	L. Gaido/INFN
0.3	23.6.2010	Added JRA1 plans (chapter 4)	L. Gaido/INFN, D.Cesini/INFN
0.4	25.6.2010	Comments by Gilles Mathieu and Cyril L'Orphelin	L. Gaido/INFN
1.0	28.6.2010	Complete draft, ready for internal review	L. Gaido/INFN
2.0	28.6.2010	New version after internal review	L. Gaido/INFN
3.0	9.7.2010	New version after G. Misurelli's review	L. Gaido/INFN
4.0	10.7.2010	New version after M. David's review	D. Cesini/INFN
5.0	20.7.2010	New version after S. Newhouse's review	P. Veronesi/INFN

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/results/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

Regionalization of the operational tools has been one of the important activities of the EGEE-III project. A large effort has been devoted to this activity during the second year of the project with the goal of providing either a regionalized instance of the tool or a regionalized view on the central instance.

However the development and testing activities took a longer time than expected so this work is not yet concluded. Regionalization of the tools is therefore an activity that is continuing within EGI's JRA1.

For each tool this document describes the status at the end of EGEE-III and the plans within JRA1.

TABLE OF CONTENTS

1 REGIONALIZATION IN EGEE-III

The EGEE-III plan about regionalization was to have three tools fully regionalized by the end of the project: the helpdesk, the Nagios-based monitoring and GSTAT [R1]. The helpdesk, based on GGUS [R2], was already a fully distributed system at the beginning of EGEE-III. During the EGEE-III project GGUS has undertaken a transition to replace e-mail-based interfaces with Web services for both import and export operations. The SAM-based monitoring infrastructure has been phased out in favour of Nagios. The GSTAT tool is presently ready for deployment by those regions willing to test regional/local BDIIs. Currently all regions operate a full project-level top-level BDII, for which a central GSTAT instance at CERN with geographical failover (at INFN CNAF, Italy) is considered to be sufficient.

The tools which are still based on a central deployment model, and will remain as such for some time during EGI-InSPIRE, are the following.

- *Messaging infrastructure.* It is based on a geographically distributed network of message brokers operated at CERN, in Croatia and Greece;
- *GOCDB [R3]*. During EGEE-III development efforts were undertaken to produce a version of the software that could also be deployed by regions. At the beginning of EGI-InSPIRE a beta package of the software is ready for testing, and United Kingdom installed an experimental instance of the service.
- *Regional dashboard*. Currently the C-COD dashboard [R4] is hosted centrally for the project operations, and regions have their regionalized views (ROD) hosted centrally. Both tools are available from the operations portal since January 2009. A beta package of the regionalized ROD dashboard software is ready for testing and United Kingdom already installed an experimental instance of the service. A demo was organized during COD-22 in January 2010 [R5].
- *Accounting portal*. Code is not ready, but a beta package for installation in the region is under preparation.
- *Accounting repository*. Code is not ready. Development efforts are part of the EGI-InSPIRE project programme of work.
- *Availability Calculation Engine (ACE)*. Will be deployed centrally for some time.

Several new National Grid Initiatives presented their regionalization plans during the Operations Tool Advisory Group meeting that was held in Amsterdam on March 03 2010 [R6]. Table provides an overall picture of the plans defined at the end of EGEE-III project. Cells with bold text indicate that the NGI wishes a central deployment of the tool with a regional “view”, providing access to a restricted subset of information relevant to the NGI scope. On the other hand, cells with courier text indicate that the NGI is planning to deploy a fully regionalized instance of the tool running under their own administrative control. Finally, cells with plain text indicate that the regionalized tool is “nice to have”, but not at top priority in the deployment plans.

Table 1: Regionalization plans of some nascent National Grid Initiatives.

NGI	Reg. nagios server	Regionalized Helpdesk	NGI Accounting repository	Regionalized Accounting portal	Regionalized Dashboard	Regionaliized GOCDB
DE, FI, NO, SE	*Needed*	**GGUS NGI view**	*Needed (SGAS)*	*Needed but using own solution*	Not needed	Not needed, will use central inst.
France	*Needed*	Nice to have/ Need evaluation of service cost of the GGUS NGI view	Nice to have: needed evaluation within French sites	**Not Needed - NGI view from central portal sufficient**	**NGI view from central portal**	*Needed (for integration with regional grids*
Germany	*Needed*	**GGUS NGI view (Germany needs a standalone helpdesk)**	*Needed, technology to be decided*	*Needed*	*Needed*	*Needed*
Greece	*Needed*	**GGUS NGI view if no cost**	*Needed (testing DGAS)*	Nice to have, available for testing	**NGI view at the beginning, regionalized when stable**	*Will use own solution*
Ireland	Needed	Own helpdesk, no NGI view	Needed, own solution based on RGMA	Not needed	TBD	Needed for integration of local sites
Italy	*Needed*	*Needed in the country, no NGI view*	*Needed (DGAS)*	*Needed (HLRmon)*	Very much needed	*Needed for integration with regional grids*
Poland	*Needed*	*Needed in the country no NGI view*	*Needed (BALTIC project)*	*Testing own solution, nice to have*	*Needed*	Not needed
Portugal	*Needed*	**interested in NGI view, if no cost**	*Needed*	*Needed*	*Needed and ongoing testing*	*Needed*
Spain	*Needed*	**interested in NGI view, if no cost**	*Needed*	*Needed*	*Needed*	*Needed*
United Kingdom	*Needed*	**interested in NGI view, if no cost**	Not needed	Not needed	**NGI view from central portal**	*Needed (testing it)*

As Table 1 shows, among the tools still to be regionalized, the deployment of an own accounting infrastructure (repository and portal) and of a regional GOCDB was a priority for most of the NGIs involved in the poll. As to the regional helpdesk and the regional dashboard, a regional view from a centrally deployed tool is considered suitable for several NGIs

The following sections describe the status of the regionalization of each operational tool at the end of the EGEE-III project. Most information was also reported in the EGEE deliverable DSA1.2.2 “Assessment of Production Service Status” [R7].

1.1 Monitoring

At the end of EGEE-III the monitoring infrastructure (regional Nagios servers [R8] and the corresponding regional MyEGEE portals [R9]) had almost completed the transition to a distributed infrastructure. This means that a Nagios instance, including a MyEGEE portal, is active in each region.

The Nagios-based system has gradually replaced the Service Availability Monitoring (SAM) infrastructure, although both systems coexisted for a while in order to validate the Nagios results.

In the last months of the EGEE project two monitoring environments have been operated: a production one and a test one.

- **Production infrastructure**. CERN hosted a central project-level Nagios and a set of regional Nagios server instances simulating the network of distributed regional Nagios servers. This set of Nagios servers constitutes the backbone of the monitoring infrastructure and is already used in production. The main consumer is the new operations dashboard, which generates notifications, and triggers alarms in case of failing critical tests. This network of Nagios servers is complemented by a set of MyEGEE portal instances, one per region, which are used by Grid operators (ROCs and site managers) to browse Nagios alarms.

  Although the availability/reliability monthly reports are still based on SAM test results, starting from March 2010 test reports are also generated in addition to the SAM ones for comparison. Since they have successfully been validated SAM are going to be phased out (English review of the last sentence).

- **Test infrastructure**. Each region currently deploys both a regionalized Nagios server and a regionalized MyEGEE portal. These set of regionalized servers is currently under validation. Once validated, the regionalized Nagios servers and the corresponding portals will gradually replace the corresponding production instances deployed at CERN. At the end of this transition, the "central" regional tool instances will be phased out. The final result of this transition will be a fully distributed monitoring infrastructure as illustrated in Figure 1.

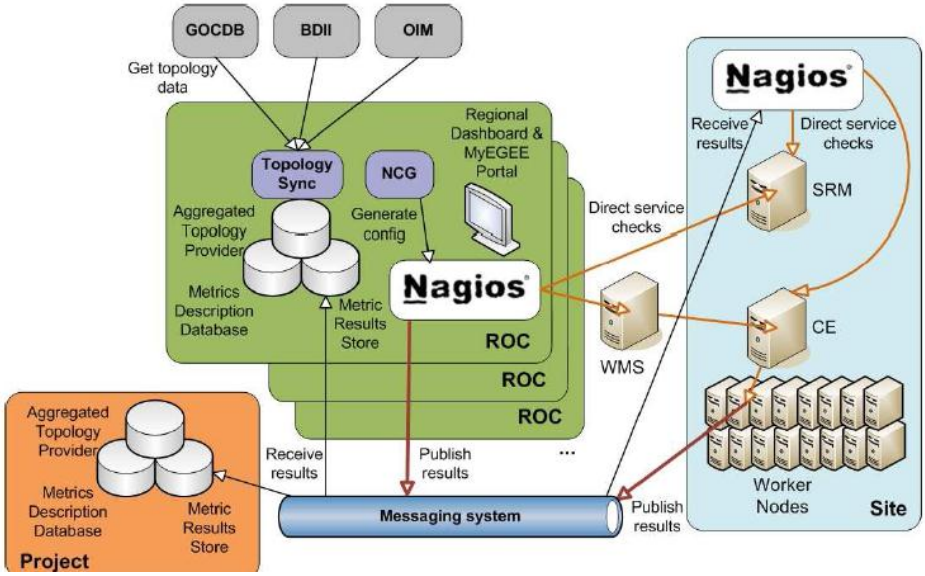

Figure 1. Configuration of the fully regionalized Nagios-based monitoring infrastructure.

A failover configuration for Nagios has to be defined.

Validation of the regional Nagios service is performed by running a dedicated Nagios server that checks:

1. the overall number of hosts and service nodes unavailable at sites in the region,
2. the functionality of Nagios processes on the server and
3. the responsiveness of the Nagios web interface.

The production messaging infrastructure is based on a network of message brokers deployed at CERN, in Croatia and Greece. Failover of the message broker network is natively implemented in ActiveMQ [R10]. If broker instances are temporarily unavailable, the message workflow is automatically redirected to the active ones to exclude any failure point.

1.2 Operations portal and dashboard

For the **operations portal** [R11] failover was implemented geographically by replicating the PHP part and the web service, as well as the Oracle backend at INFN CNAF (Italy). To switch from the primary instance – hosted in IN2P3 (France) – to the secondary, a manual intervention was requested at the DNS level. The synchronization between secondary and primary instances required about one hour downtime to complete the setup. This geographical failover mechanism will be maintained until a new setup within EGI will be completed.

The **operations dashboard** [R4] has been totally developed during the second year of the EGEE-III project. The tool was ported to Symfony [R12] and the backend – collecting results from the monitoring infrastructure and originally based on SAM – was ported to Nagios and the messaging system, as illustrated in Figure 2.

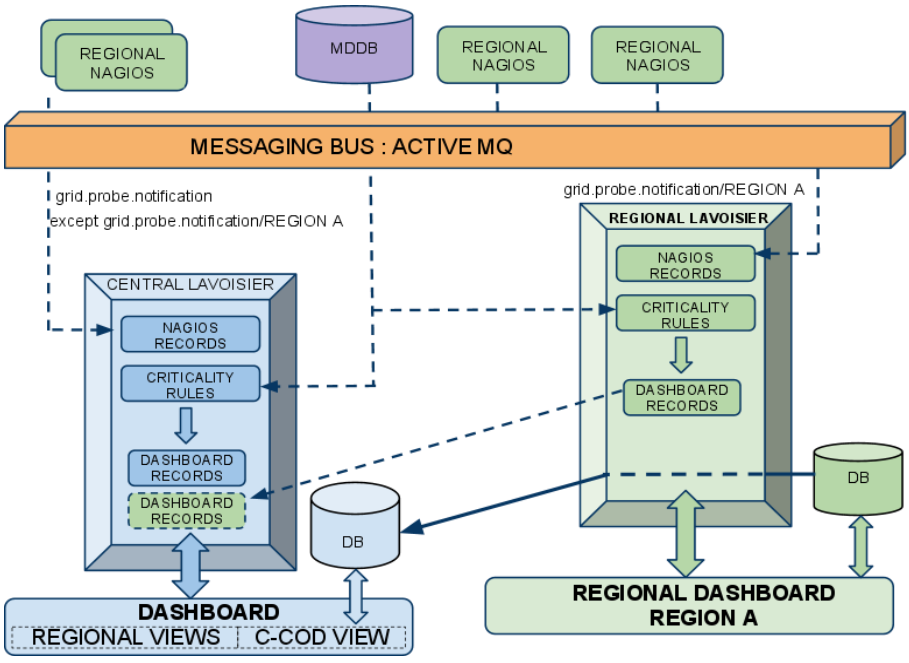

Figure 2: Architecture of the new dashboard based on Symfony and adapted to the Nagios monitoring framework and to the messaging communication system.

The new dashboard has been subject to several validation campaigns and finally put into production at the beginning of March 2010. The new dashboard relies on local failover. In particular:

- *lavoisier*: a spare machine is run in stand-by mode to replace the production one as needed;
- *database*: a local DB can be restored from the backup in 30 minutes.

A MySQL-based high availability configuration is foreseen for local automatic replication of the database (the MySQL database will be integrated into a cluster), whereas the web front-end is already implemented as a cluster.

1.3 Accounting portal

The EGEE Accounting Portal [R13] is the graphical front-end to the EGEE accounting infrastructure. Accounting statistics are available through the Accounting Portal for the analysis by different grid users, VO administrators and site administrators.

A test portal instance runs in parallel with the production one; both are hosted by Xen virtual machines running on two different physical servers. The production instance runs a stable software release, whereas the test instance is used for beta testing of new software releases. This test instance can also be used in case of failures of the primary one, as usually the test version of the portal is sufficiently stable for production.

In case of major failures of one or both machines several recovery strategies are possible depending on the technical issue encountered.

- Missing files can be restored from the backup.
- A virtual machine can be easily migrated to a different physical host, as virtual machine images are periodically saved.
- Software releases can be easily re-installed from the local subversion.
- The local usage record database is periodically generated from the central accounting ones, and a synchronization script can be manually run to download fresh accounting data from the remote databases. A few additional extra datasets are usually stored in static local databases, and can be restored from various locations (these data are equally present on both the production and test service instance, and can be also restored from the general site backup utility).

A recent version of the tool has been released in the last months of the EGEE project, including new features such as:

- the creation of per-country charts showing the percentage of jobs run by international Grid users from other countries;
- views showing the percentage of sites publishing User DNs;
- the possibility to select any VO of choice of the user in the overall view;
- the possibility to choose the HEP Spec benchmark instead of Spec Int 2000.

The existing code base has then been frozen.

1.4 Accounting Repository

APEL (Accounting Processor for Event Logs) [R14] is a CPU usage accounting tool. APEL is deployed in many EGEE regions for the collection of accounting information from individual sites and other Grid

accounting domains relying on different accounting infrastructure, such as OSG – based on GRATIA [R15], the Italian Region – based on the Distributed Grid Accounting System (DGAS) [R16] and the federation of the Nordic countries – based on SGAS [R17], as illustrated in Figure 3.

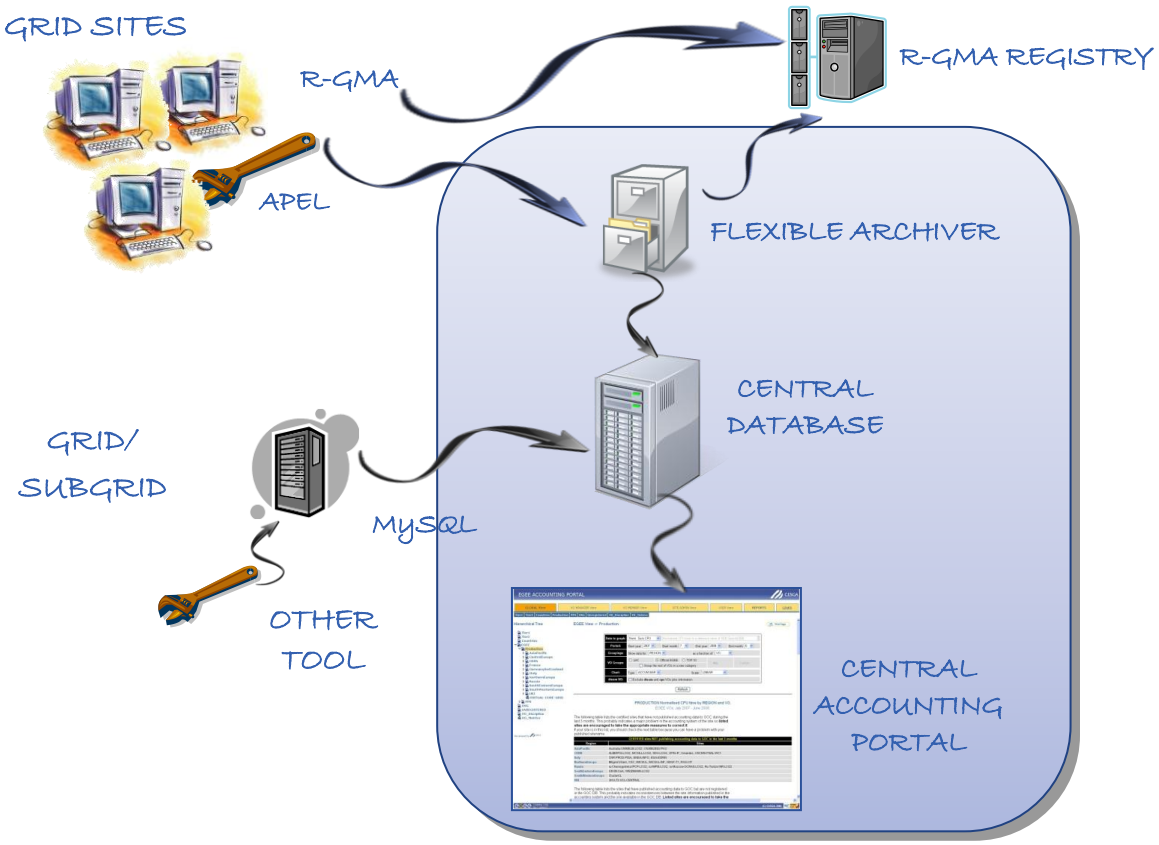

Figure 3: Current R-GMA-based APEL architecture.

As a log processing application, APEL interprets logs of Grid gatekeeper (e.g. globus) and batch system logs (e.g. PBS, LSF, SGE and Condor) to produce CPU job accounting records identified with Grid identities. APEL publishes accounting records into a centralized repository at a Grid Operations Centre (GOC) operated at RAL (United Kingdom), for access from the EGEE accounting portal. The functions of log files parsing, record generation and publication are implemented by the *APEL Parser*, *APEL Core*, and *APEL Publisher* component respectively. The APEL parser and publisher constitute the APEL client installed at the sites (on the Compute Element and on the MON box node respectively).

In the original architecture the production system used R-GMA (Relational Grid Monitoring Architecture) as the transport mechanism for moving accounting records generated on each Grid client site to a centralized repository at a GOC. R-GMA Primary Producers publish records from each Grid site and a Secondary Producer aggregates records into a centralized repository. A general topic publication and subscription messaging model enables distributed components in a system to

publish and subscribe messages to/from a well defined topic that can be viewed as a virtual destination and source of message.
The replacement of R-GMA by a new transport mechanism (ActiveMQ), expected by the end of EGEE-III, has happened, although with a small delay: AMQ server entered production in May and the AMQ client was rolled out as part of gLite at early June and both architectures are now in production.
APEL is still in a transition phase because a large number of sites still use the RGMA channel.

1.5 GOCDB

The GOCDB architecture comprises a front-end (Web services and graphical user interface) and a backend (Oracle database) offering persistent storage of the Grid configuration data, as illustrated in Figure 4.

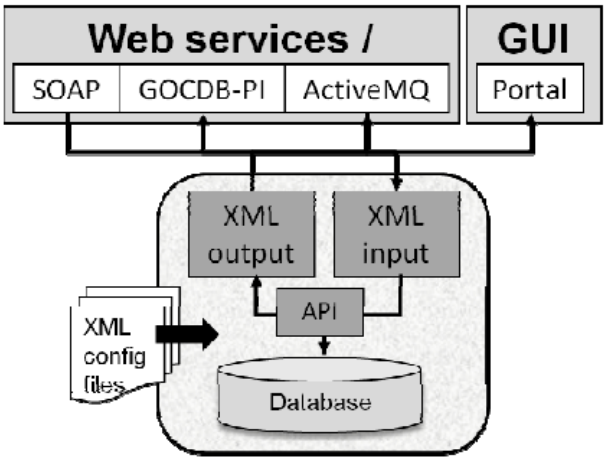

Figure 4: GOCDB architecture.

The current GOCDB implementation (GOCDB3) is going to be replaced by a new version (GOCDB4), mainly developed to enhance the internal data representation and database structure, but for some time the two instances of the tool will still be run in parallel. GOCDB3 acts as the "central input system" (where users update information), whereas the GOCDB4 input system is just available in test mode at the moment. GOCDB3 data are then pushed into GOCDB4 every 10 minutes.

The full GOCDB4 deployment will be possible with the replacement of GOCDB3 by the GOCDB4 input system, and has the following prerequisites:

- GOCDB4 input system is fully implemented and has passed all tests;
- any third-party tools connecting to GOCDB have migrated to the GOCDB4 version of the programmatic interface where needed.

The failover architecture of GOCDB comprises: a replica portal for GOCDB3 at ITWM (Germany), a Oracle database replica hosted at INFN CNAF (Italy), and a local database replica at RAL (United Kingdom). Production data are replicated every two hours.

The replica portal is kept synchronised with the master one by regular RPM updates, while the replica databases (both local and remote at CNAF) allow for data dumps to be imported and used replacing the master database. This dump process is currently manual.

The GOCDB front-end URL is provided through a top level DNS (gridops.org) that can be switched to point to either the master or replica instance thus making any intervention for the setup of the failover configuration transparent to the users. This is of paramount importance in order to minimize the impact on third party tools relying on the GOCDB programmatic interface and requiring a static URL instance.

The general failover architecture makes possible to combine any of the two portal instances and of the three database instances interchangeably. However, the most common configurations are:

- Master portal and master database;
- Replica portal and replica database
- Master portal and local backup database.

The replica portal is usually set up to work with the replica database, this implying that the only operation needed in case of outage is to swap DNS entry to point to the replica front-end. On the other hand, using a master portal with a backup database implies importing the data to the backup first, and then to configure the master portal to use this database.

Enhanced solutions for database synchronisation will be put in place for GOCDB4 using Oracle streams.

1.6 GGUS

The EGEE helpdesk is a distributed system with central coordination. The central part (GGUS) is run on two distinct GGUS servers, one dedicated to production services and the second to training. Of the two servers constituting the *front-end*, one is active and the secondary is in standby. The primary is hosted by a physical machine, whereas the secondary is on a virtual host (Vmware ESX cluster with Dynamic Resource Sharing and High Availability and a redundant connection to a Storage Area Network). At the time of writing switching from primary to secondary still requires manual intervention. The design of an automatic failover mechanism is part of the future development plan.

The underlying SAN infrastructure at KIT is distributed at two different locations, with independent power supply. GGUS currently relies on RedHat Enterprise distribution v.5 and Remedy has been recently updated to version 7.1. Network connectivity is redundant and obtained from two different network providers.

The database at the *backend* is divided into three instances dedicated to production, development and training. It has recently been updated and moved to Oracle 10g Real Application Cluster (RAC) with two nodes and DataGuard at two separate locations for replication.

GGUS servers are under monitoring by the local Nagios instance and an alarm is triggered in case of failures outside office hours. On-call staff of the “Grid services and GGUS team” are responsible for interventions outside office hours. Incidents concerning the functionality of GGUS are of service level 1, this meaning that they must be dealt with instantaneously.

1.7 Metrics portal

The main objective of the metrics portal [R18] is to have a set of metrics that can help to measure project performance and keep track of its evolution. The development of this tools within EGEE has been driven by the Operation Automations Team (OAT). The portal automatically collects data from different sources (e.g. GOCDB [R3], GGUS [R2], GridView [R19], etc.) by means of various connectors which depend on the data provider and displays them on the portal. The connectors translate the information gathered from different producers and store it into a local database. The portal is an easy and fast way to have all these information reported by means of a common interface.

2 OPERATIONAL TOOLS REGIONALIZATION IN EGI-INSPIRE ACCORDING TO JRA1 PLANS

According to the EGI-InSPIRE proposal the operational tools development will be carried out by a number of Operations Product Teams that will be responsible for the maintenance and development work of each operations tool. Each Operations Product Team can be geographically distributed, with different NGIs contributing to the development of the work. Technical collaboration between the operations Product Teams and the NGIs will take place through the Operations Forum. The product teams are coordinated by and their long-term road map is determined through the Operational Tools Advisory Group (OTAG). JRA1 acts as a glue to monitor and coordinate the activities carried out by the different actors, follows up the tool-maintenance roadmaps and ensures that software releases are successfully deployed in production.

OTAG's main goal is to streamline the requirement gathering process which concerns the various operations Product Teams, coordinated under the umbrella of the Operations Automation Team, and the tool customers (regional operations centres, on-duty operations personnel, site managers and Grid end-users). The OTAG mandate is also to manage the development and evolution of operational tools in response to new scenarios and feedback from its users. New requirements are collected and prioritised, and ongoing development, testing and release activities are reported. OTAG provides a forum to discuss the future evolution of the operations tools and to agree tool roadmaps that meet the expressed needs of the EGI community. The first OTAG meeting was held during the last phase of the EGEE-III project (in Jan 2010) and the group meets on a monthly basis since then, reporting directly to OAT.

In the following sections for each tool the maintenance and development plans, that are the result of the OTAG activity and of the discussions at the JRA1 kickoff meeting [R20], held in Amsterdam at the beginning of June, are reported. The road map for each, as defined by JRA1, will be fully described in milestone MS703 (Operation Tool regionalization work plan) that is currently being prepared.

2.1 Monitoring

The regionalization of the SAM system was almost completed during EGEEIII as described in section 3.1. During the first month of the EGI project a first update to the regional NAGIOS has been released in order to fix a number of bugs for all the web components. It also provides new probes (for glexec and MPI) and a Nagios Grid Information Provider.
The validation process of the regional instances is still ongoing and the progress is tracked on a wiki page [R21].
At the moment of writing this is the NAGIOS boxes validation status:

- 14 ROCs/NGIs successfully validated and completed the migration (Central Europe, CERN, France, Germany&Switzerland, Italy, Northern Europe, South East,, UKI, Latin America, IGALC, NGI_PL, Greece, Serbia and Slovakia).
- 2 ROCs/NGIs have already validated but not migrated (South West Europe and Slovenia))
- 3 ROCs not started yet (Asia Pacific, Canada and Russia)

On June 1st the central NAGIOS instance has been switched off for those NGIs/regions that completed the validation and availability/reliability statistics are calculated using the data from their regional boxes.

Future developments of the regional boxes concern the myEGI portal, that will be bundled with the NAGIOS itself. It will provide specific NGI views and will be the portal for information about availability, reliability and service status of NGI resources; 'GridMap style' views will be added later on.

JRA1 development plans include also a Metric Description Data Base interface to allow the web editing of profiles and the Aggregated Topology Provider (ATP). The Nagios central instance will gather topology information from various sources (for example top-level BDII and GOCDB) about hosts and the respective services supplied, and this information will be synchronized with the ATP instances deployed regionally as illustrated in Figure 1.

It is worth highlighting that a clear and well documented procedure to insert new probes (with potentially new metrics) into the Nagios monitoring infrastructure as well as a straightforward procedure to distribute it to the NGIs is considered very important.

2.2 *Operations portal and dashboard*

The Operation portal is currently only a central service. However during the first three years of the project the CIC portal features will be ported to the Operations Portal [R11] and its functionalities that are meaningful at a NGI Level will also be provided into the regionalized portal package. The features and plans about the operations portal are described in the milestone MS701 describing the roadmap for the CIC Operations portal [R22].

Currently three instances of the regionalized package have been deployed: one (NGI_GRNET) has already been validated and the remaining two (NGI_IBERGRID and NGI_CZ) are in the validation process at the time of writing.

The central service will be maintained for those NGIs that will not deploy a regionalized instance when it will be available. Development plans include the porting to symphony, the integration with other portal framework (in particular with the GOCDB), the adoption of portlet/widget technologies in order to make some of its pages pluggable into other scientific gateways. It will evolve towards a service oriented model and information will be made available through a series of protocols and channels that will make them easily usable by the end users (local language, mobile browsers, etc.).

At the end of EGEE-III the operations dashboard already contained a series of regional views according to the project ROC structure and it will gradually evolve towards a regionalized model. The first regional package for the dashboard (to be deployed at the NGI level) was released on June 8th.

The central application now implements a fully regionalized model with a synchronization system to exchange information with the upcoming regional instances, shown in Figure 5.

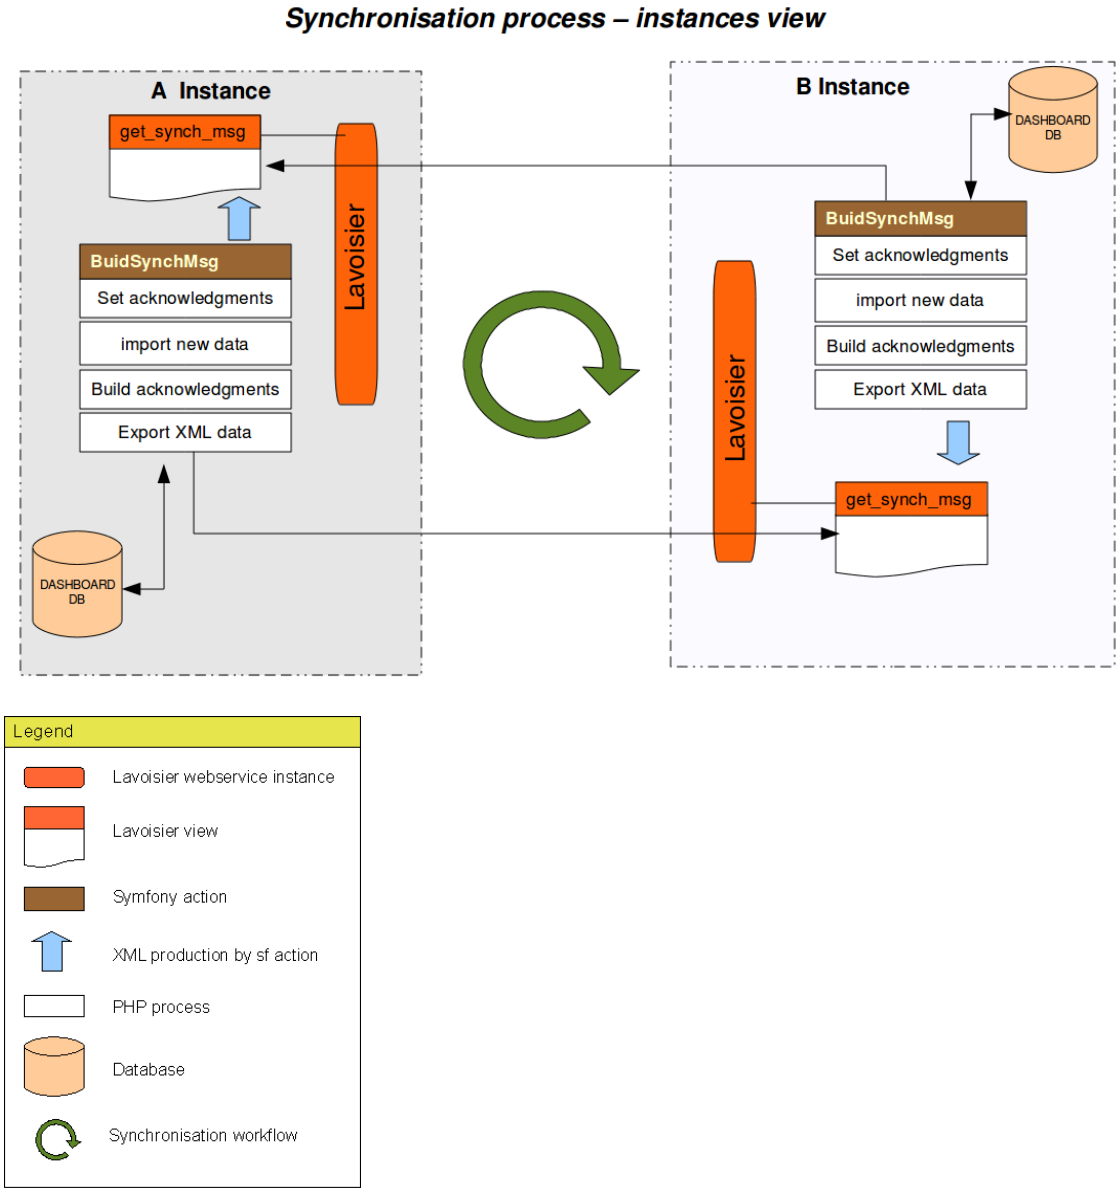

Figure 5: The synchronization system for the regional dashboard

The package is distributed in two modules: the Lavoiser web service (installed through rpm) and the php pages (to be checked out directly from SVN).

2.3 Accounting portal

Since a regionalized flexible archiver is not ready for regional installation at the time of writing, the central accounting portal needs to be extended to support regional deployments of NGIs accounting systems and national deployment of the accounting repository. Two complementary scenarios will be available: regional accounting portal released for those NGIs that wish to install it and a new NGI view on the central instance for those not wanting the additional effort of maintaining their regional instance.

2.4 Accounting repository

The regionalisation of the APEL server has been somewhat delayed with respect to the original plans. Therefore the tool is not distributable in its current state and this is not likely to happen until the second year of the project.

2.5 GOCDB

The project will provide an instance of the GOCDB service that could be deployed nationally and information federated into a central instance. New developments concern GOCDB4 in order to allow an easy distribution and regionalisation of the tool. There will be a GOCDB4 module installed by NGIs to store their topology information ("GOCDB4 Regional Module"), a central system ("GOCDB4 Central Instance") to collect all EGEE/EGI data from the various regional instances, a central input system ("GOCDB4 Input System") to provide a means to access in write mode the central instance to those NGIs that do not run their own regional GOCDB4. The GOCDB4 central instance was released in November 2009 in read only mode for the users, data insertion is done through a synchronization system from GOCDB3 as depicted in Figure 6.

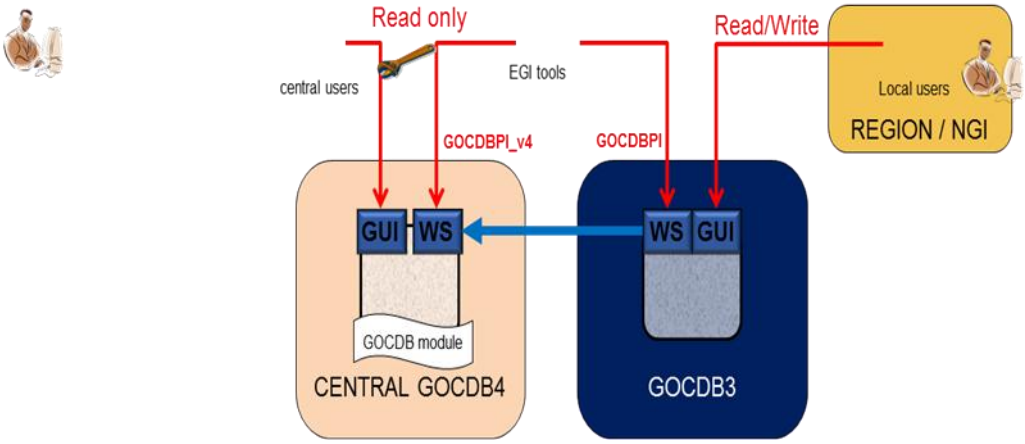

Figure 6. Syncronization between GOCDB3 and GOCDB4

The GOCDB4 full release will consist in the replacement of the GOCDB3 with the GOCDB4 input system as in Figure 7.

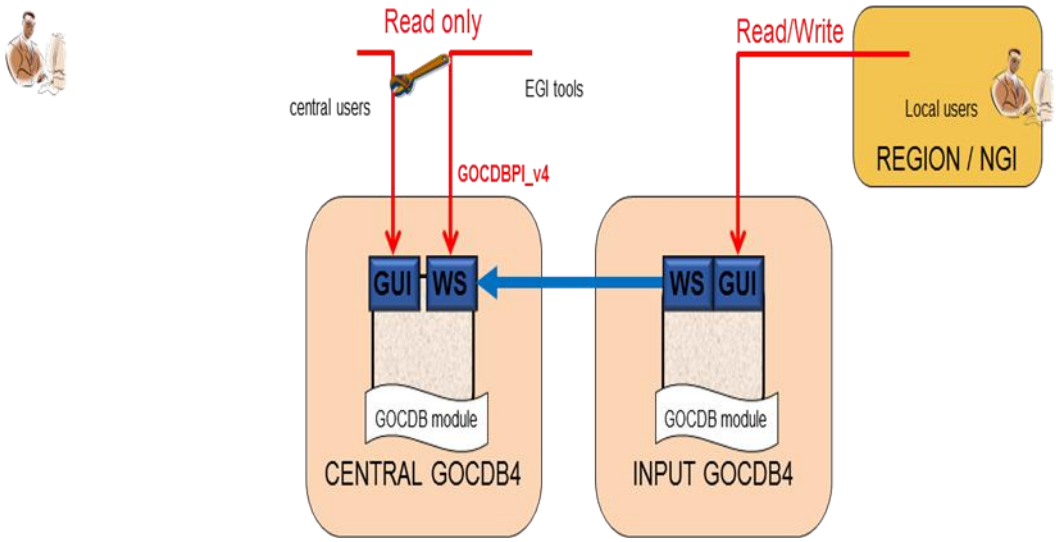

Figure 7: Final GOCDB4 topology

Two steps will be done starting from mid June in order to complete the migration:

- decommissioning of GOCDB3 programmatic interface to be done that all the dependencies compatibility issues of the other operation tools are solved.
- replacing GOCDB3 with GOCDB4 input system

A test version of the GODB4 regional module was release during EGEE-III not yet used in production. The production quality regional GOCDB4 is foreseen to be released on mid July. When regional module will be available, the NGIS will have the following options:

- use their own GOCDB, automatically synchronized with the central instance
- use central instance populated through the GOCDB4 input system as described above
- use any other tool with functions similar to the ones of the regional module

2.6 GGUS

The central Helpdesk system (GGUS) is not foreseen to be deployed at the NGI level. Options for the NGIs are:

- Direct use of GGUS (small, upstarting NGIs)
- Regional View of GGUS ("EGI-only" NGIs)
- Regional Helpdesk System (large, heterogeneous NGIS)

The regionalized view is currently under testing by the NGI-DE. The third scenario has already been implemented by most NGIs during the EGEE projects by means of various different ticketing systems. However some of them have recently shown an interest in using a popular tool (RT); testing activities are ongoing.

2.7 Metrics Portal

No regional model is foreseen; the portal will run on a central install only.

3 CONCLUSIONS

JRA1 has the primary goal of coordinating the operational tools development carried out by various groups. The guidelines are being defined with the fundamental support of the OTAG advisory group.

The validation process of the Nagios regional instances, to replace the old SAM system for the availability/reliability metrics, has not been concluded yet although 14 regions/NGIs are already using it in production (and two more completed the validation).

The first regionalized package of the operations portal and dashboard is being tested by three NGIs; its development will continue and improvements are foreseen although the central instance with a regionalized view is not going to be phased out.

The accounting portal and repository are not ready to be regionally deployed; development toward regionalization is planned but a first prototype will not be ready before the second year of the project.

The 'production-quality' regional module of GOCDB4 is going to be released very soon.

No regionalization is currently envisaged for GGUS and for the Metrics Portal.

4 REFERENCES

[R1]	*The GStat portal* (http://goc.grid.sinica.edu.tw/gstat//index.html)
[R2]	*Global Grid User Support (GGUS) Plan*, *EGEE-III Deliverable DSA1.1, August 2008* (https://edms.cern.ch/document/931935)
[R3]	Mathieu, G.; Casson, J.; *GOCDB4, a New Architecture for the European Grid Infrastructure*; to appear in the Proc. of the Int. Symposium on Grid Computing 2010, Mar 2010, Academia Sinica, Taipei, Taiwan
[R4]	*The Operations dashboard* (https://operations-portal.in2p3.fr/index.php)
[R5]	*COD-22 Face-to-Face Meeting*, Jan 2010 (http://indico.cern.ch/conferenceTimeTable.py?confId=68900#20100126)
[R6]	*Operational Tools Advisory Group meeting, Amsterdam, March 03 2010* (https://www.egi.eu/indico/conferenceDisplay.py?confId=7)
[R7]	*EGEE-III DSA1.2.2: Assessment of production service status* (https://edms.cern.ch/document/1060571)
[R8]	*The Nagios project web* (https://ops-monitor.cern.ch/nagios)
[R9]	*MyEGEE Portal* (https://twiki.cern.ch/twiki/bin/view/EGEE/MyEGEE#MyEGEE_Portal)
[R10]	*The Apache ActiveMQ messaging system* (http://activemq.apache.org)
[R11]	*The CIC Operations Portal* (https://cic.gridops.org)
[R12]	*The Symfony PHP web framework* (http://www.symfony-project.org)
[R13]	*EGEE Accounting Portal* (http://www3.egee.cesga.es)
[R14]	Jiang, M.; Del Cano Novales, C.; Mathieu, G. et alt; *An APEL Tool Based CPU Usage Accounting Infrastructure for Large Scale Computing Grids*, *in the Proc. of ISGC 2010,* Mar 2010, Academia Sinica, Taipei, Taiwan
[R15]	*The OSG Gratia accounting system* (https://twiki.grid.iu.edu/bin/view/Accounting/WebHome)
[R16]	*The DGAS accounting system* (http://www.to.infn.it/dgas)
[R17]	*The SGAS accounting system* (http://www.sgas.se)
[R18]	*The EGEE Metrics Portal roadmap* (http://www3.egee.cesga.es/gridsite/accounting/CESGA/links/EGEE-III-SA1-Metrics_Portal_Roadmap_and_Requirements-v3.1.xls)

[R19]	Sonvane, D.; Kalmady, R.; Chand, P.: *GridView Service Availability Computation, August 2008* (https://twiki.cern.ch/twiki/pub/LCG/GridView/Gridview_Service_Availability_Computation.pdf)
[R20]	*The JRA1 kickoff meeting, Amsterdam, June 2010* (https://www.egi.eu/indico/conferenceDisplay.py?confId=46)
[R21]	*Validation of Nagios installations wiki* (https://twiki.cern.ch/twiki/bin/view/EGEE/ExternalROCNagios*)*
[R22]	*Milestone MS701: roadmap for the CIC Operations Portal* (https://documents.egi.eu/public/ShowDocument?docid=39)