



EGI-InSPIRE

ESTABLISHING THE OPERATIONAL TOOL PRODUCT TEAMS

EU MILESTONE: MS702

Document identifier:	EGI-MS702-final
Date:	05/11/2010
Activity:	JRA1
Lead Partner:	INFN
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/52

Abstract

This document presents the operational tool product teams working in the WP7 (JRA1) activity. Details on their development and testing infrastructure are also given after a brief introduction to the developed tools. The common release procedure and the relationship with external bodies in order to get new requirements are described.



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	Daniele Cesini	INFN	29/06/2010
Reviewed by	Moderator: Gergely Sipos Reviewers: Steven Newhouse	EGI.eu	01/07/2010
Approved by	AMB & PMB		26/07/2010

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1.0	29/06/2010	First draft	Cesini
2.0	01/07/2010	Received moderator review	Sipos
3.0	01/07/2010	Received project director review	Newhouse
4.0	19/07/2010	New document produced and uploaded	Cesini
5.0	22/07/2010	Received second review from moderator	Sipos
6.0	22/07/2010	Created new version (5.0)	Cesini
7.0	26/07/2010	Uploaded version 5.0 on document repository	Cesini

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.



V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:

<https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY

A complete project glossary is provided at the following page: <http://www.egi.eu/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

This document presents the status of the geographically distributed operational tools product teams participating to the WP7 (JRA1) at the beginning of the project. A report is given about their location, manpower, development infrastructures, facilities used, testing procedures, communication channels, etc. Each product team has a specific and independent infrastructure, mostly inherited from the past EGEE projects and it was decided not to change this structure, having a central coordination and a common release procedure that is described at the end of the document. The activity management will assure proper internal and external communication among product teams and with other project bodies or advisory groups through periodic phone conferences and face to face meetings.

The document is composed of eight sections. In section 3 an overview of the JRA1 activity according to the project description of work is given. Section 4 presents the development infrastructure and the characteristics of each product team. Section 5 is about the interactions of the product teams, both internally to the activity and externally with other project bodies. Section 6 presents the release procedure common to all product teams and their responsibilities. Section 7 is about the future of product teams and their development infrastructures.



TABLE OF CONTENTS

1	INTRODUCTION: JRA1 ACTIVITY OVERVIEW	7
2	THE PRODUCT TEAMS AND THEIR DEVELOPMENT INFRASTRUCTURE	8
2.1	The operational tools product teams.....	8
2.1.1	Operation Portal PT and development infrastructure.....	10
2.1.2	EGI Helpdesk PT and development infrastructure.....	11
2.1.3	Grid Configuration Database (GOCDB) PT and development infrastructure.....	12
2.1.4	Accounting Repository PT and development infrastructure	12
2.1.5	Accounting Portal PT and development infrastructure.....	13
2.1.6	Service Availability Monitoring PT and development infrastructure.....	14
2.1.7	Message Bus Configuration PT and development infrastructure	14
2.1.8	Metrics portal PT and development infrastructure.....	15
3	INTERACTIONS OF THE PRODUCT TEAMS.....	17
4	OUTPUT OF THE PRODUCT TEAMS	18
4.1	Strengthening the test phase and monitoring the software quality.....	19
5	FUTURE OF THE PRODUCT TEAMS	20
6	CONCLUSIONS.....	21
7	REFERENCES	22



1 INTRODUCTION: JRA1 ACTIVITY OVERVIEW

The JRA1 activity is responsible for the development and continuous evolution of the operational tools used by the European Grid Infrastructure. Future developments will follow the evolution of the infrastructure that is expected to encompass additional resource types, such as desktop grids, cloud and high capacity computing resources. The work started during the EGEE-III project to allow for a fully distributed or regionalized model of the operational tools should be completed during the early stages of the activity. The evolution of the accounting system towards an economic and business model is also an objective of the activity. The activity is composed by five tasks:

- TJRA1.1 – Activity management: is the coordination of the development work, including the definition of the release roadmaps, the representation of the activity within the projects bodies and the reporting on the status of the activity. Runs for 4 years.
- TJRA1.2 – Maintenance and development of the deployed operational tools: is the underlying bug fixing and development work for the operational tools. A general description of the tools will be given in section 3. This task runs for 4 years.
- TJRA1.3 – National Deployment Models: is the development needed to support the deployment at a regional (or NGIs) level. This activity, started during the EGEE-III project, should be completed by the end of the first year of the project.
- TJRA1.4 – Accounting for different resource types: is the work needed to evolve the EGI accounting system in order to encompass the different resource types that will be included into the production infrastructure and to support an economic model needed for the self sustainability of the EGI Grid. This is a 3 years task starting from the second one.
- TJRA1.5 – Integrated Operation Portal: the development needed for a restructuring of the Operation Portal and its harmonization with other portal frameworks such as the GOCD. This task spans over the first three years of the project.

The relationship between the JRA1 activity and other project bodies is described in section 5.



2 THE PRODUCT TEAMS AND THEIR DEVELOPMENT INFRASTRUCTURE

In this section the product teams (PT) participating to the activity are introduced, a brief description of their developed software and details about their development infrastructure are also given. Table 2 at the end of the section summarises some of the infrastructures details.

2.1 *The operational tools product teams*

A product team is a group of developers responsible for developing, testing, maintaining and supporting software and middleware used by the project. The operational tools product teams (OTPTs) develop tools that are fundamental during the day by day operation of the production Grid and are widely used by all the actors involved in running and supporting the infrastructure, i.e. sitemanagers, NGI managers, support teams, etc. OPTPs are the only software providers internal to the EGI-InSPIRE project, all other pieces of software deployed and used by the project come from third parties or other projects. So OTPTs are the only type of PTs that operate inside EGI-InSPIRE working in a single project with their customers – and this is a similarity with middleware development teams of the EGEE projects – however, the development and operation of EGI Operation tools now take place in two different work packages, which is a shift from EGEE, where these happened inside a single work package, EGEE SA1.

The OTPTs can be geographically distributed with more than one partner contributing to it. Each OTPT is independent and uses its own development infrastructure but the development efforts are coordinated centrally by the activity management and the software roadmap for each tool is discussed and agreed within the activity through periodic meetings and are driven by requirements coming into JRA1 from external bodies such as the User Services Advisory Group (USAG) or the Operational Tools Advisory Group (OTAG).

OTPTs will collaborate with an external software provider (e.g. EMI project) when the operational tools interface with one of their software components. Any new features in the software component required to support this integration will be approved through the Technical Coordination Board (TCB).

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

- **Operation Portal:** (also known as the CIC Portal) a single access point to many operational information and operational tools. Used by all the project actors, i.e. end-users, sitemanagers, NGI managers, support teams, VO managers. The product team is run by the CNRS partner located in Lyon, France. The portal has a single central installation [2] but a regional package is being developed and deployed on the interested NGIs.
- **EGI Helpdesk:** the main support access point for the project. Primary exploited by Grid end-users and by support teams, but also by sitemanagers who need support to solve middleware issues. The product team is run by the KIT partner located in Karlsruhe, Germany. The portal has a single central installation [7].
- **Grid Configuration Database (GOCDB):** contains general information about the sites participating to the production Grid. Accessed by all the project actors (end-users, sitemanagers, NGI managers, support teams, VO managers), by other tools and by third



party middleware in order to get Grid topology. The product team is run by the STFC partner and is located at RAL, Didcot, UK. The portal has a single central installation [33] but a regional package will be developed and deployed on the interested NGIs.

- **Accounting Repository:** stores information about the usage of sites within the production infrastructure. Primary accessed by other tools (eg the Accounting Portal) in order to create statistical reports. Developed by the same OTPT as GOCDB at RAL.
- **The Accounting Portal:** graphical frontend for the Accounting Repository. Used by anybody interested in accounting information. The product team is run by the partner CSIC and is located at CESGA, Santiago de Compostela, Spain. The portal has a single central installation [34] but a regional package will be developed and deployed on the interested NGIs.
- **The Service Availability Monitoring:** a monitoring framework for the resources that participate in the production infrastructure. Used by sitemangers to monitor their sites but also at a project level as data source to create availability/reliability statistics. The PT is geographically distributed across three partners: CERN (Genève, Switzerland), SRCE (Zagreb, Croatia) and AUTH (Thessaloniki, Greece). The service is now regionalized with multiple instances deployed across the EGI Grid.
- **The Metrics Portal:** collects a set of metrics from different resources to help measure project performance and keep track of project evolution by displaying historical values of the metrics in a single place. Developed by the same PT as the Accounting Portal at CESGA. The portal has a single central installation only [35].

The JRA1 activity is also responsible for providing support to the integration of the tools with the message broker network run by EGI. This is not a tool developing work in the strict sense but it is a cross tools activity carried on by a PT located at AUTH.

As anticipated in section 3 an activity task, starting in the second year of the project, will also add functionalities to the accounting system currently used within the project in order to account for more resource types. The PT and hence the development infrastructure for this task are still not well defined, but the work will be done in strong connection with the already existing accounting development infrastructures inside the EGI (section 4.1.4) and EMI projects. The effort for this activity will be spread over four partners: LUH(Hannover, Germany), INFN (Italy), SFTC (UK) and CESGA (Spain).

The map in Figure 1 shows the geographical distribution of the different PTs, details about the development and testing infrastructure are given in the next subsections.

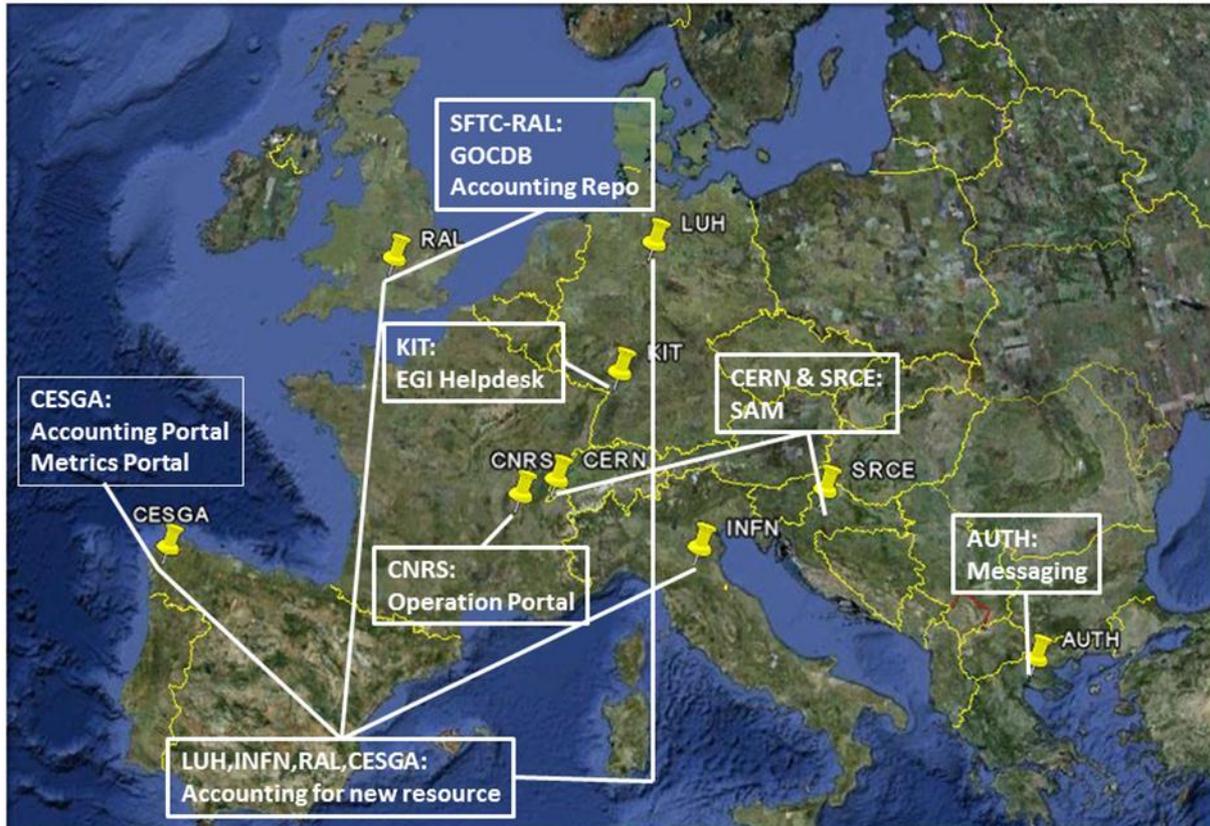


Figure 1 - Geographical distribution of Product Teams

2.1.1 Operation Portal PT and development infrastructure

The Operation Portal is an integration platform, allowing for strong interaction among existing tools with similar scope and filling up gaps wherever functionality was lacking. It also implements numerous facilities derived from requirements expressed by end-users, VO managers, NGI managers and sitemanagers. It provides information retrieved from several different distributed static and dynamic sources – databases, Grid Information System, Web Services, etc. – and gathered onto the portal. Its user tracking, failures monitoring, alarm notification and ticketing system interaction capabilities are exploited during the day by day operation of the production Grid. The portal also fosters communication between different actors through channels like the Broadcast Tool and the Downtime Notification Mechanisms, putting in place procedures to address their interactions needs.

All the key features of portal are being migrated into the Symphony [1] php web framework to improve the efficiency of the Operations Portal in term of response time, to decrease the time spent on maintenance and to increase robustness by developing independent modules. The release of the first production regionalized package is from June 8th 2010.

Further architecture and implementation details can be found on the portal itself at [2].

Description of the production and testing infrastructure:

- **Infrastructure location:** CNRS-IN2P3, Lyon, France
- **Team composition:** three people (68 PMs over 4 years)
- **Distribution of work inside the team:** two people are responsible for Symphony [1] developments and one for the web service integration and the global architecture.

- **Code repository used:** code managed through SVN [5]
- **Build system:** Maven [3] is used to build the web service and Symfony scripts for the remaining code.
- **Bug/task tracking facilities:** the internal bug/task tracking system is hosted on the IN2P3 Forge system [4].
- **Internal communication channel:** cic-information<at>in2p3.fr mailing list
- **Support communication channels:** cic-information<at>in2p3.fr mailing list and a request form on the Portal itself. In case of bugs a ticket can be opened either on GGUS (section 4.1.2) or directly on the internal bug tracking system.
- **Testing and validation procedures:** two completely independent infrastructures are in place for development and testing, but can be synchronized via SVN. Once the development is done and validated on the development infrastructure the code is propagated to the test infrastructure where the new features are first checked by the developer in charge and after by one of the other developers.
- **Distribution facilities:** software distribution is done via SVN [5].

2.1.2 EGI Helpdesk PT and development infrastructure

The EGI Helpdesk (GGUS) is the primary way to request support about any production infrastructure related problem. Trouble tickets are created upon request and are used to track the problem until a solution is reached. The Helpdesk system is not foreseen to be deployed at the NGI level, but regionalized views are currently under test.

Details on the architecture and implementation can be found on the main GGUS documentation page [7].

Description of the production and testing infrastructure:

- **Infrastructure location:** Developers and development infrastructure are located at KIT in Karlsruhe
- **Team composition:** Four developers (47 PMs over 4 years)
- **Distribution of work inside the team:** one developer is responsible for the frontend, one for the Remedy [6] integration, one for the failover and system administration, one for the regional view development.
- **Code repository used:** SVN is used for code management
- **Build system:** not built
- **Bug/task tracking facilities:** feature requests and bug tracking is done through a savannah project hosted at CERN [8]. The request tracker contains a “shopping list” that is an open forum where wishes’ and requirements are recorded and discussed by an advisory group. The decisions made by this group are then published on the Ongoing Worklist [9] which is the official plan where everybody can get information about what is going on GGUS.
- **Internal communication channel:** The GGUS mailing list (ggus-info<at>cern.ch) can be used to contact the developers – it will be moved soon to an EGI address.
- **Support communication channels:** The GGUS mailing list and a support unit in GGUS itself
- **Testing and validation procedures:** New releases are timely announced via various communication channels and the new features are moved to the GGUS training system where are intensively tested by the GGUS developers and by the regional helpdesk developers if the interfaces to other helpdesk systems are involved by the new code. Automatic test scripts for basic functionalities are also available. Manual testing is always done after each release for key features, such as the “Alarm Ticket” chain.

- **Distribution facilities:** not distributed

2.1.3 Grid Configuration Database (GOCDB) PT and development infrastructure

The Grid Configuration Database (GOCDB) contains general information about the sites participating in the production Grid: the installed middleware services, the administrators and security contacts, the scheduled and unscheduled downtimes are examples of the information stored into the GOCDB. This set of information is available through a series of web pages accessible via X509 certificates. The project will provide an instance of the GOCDB service that could be deployed nationally and federate information into a central instance.

Implementation details and architecture are available at [10] and [11].

Description of the production and testing infrastructure:

- **Infrastructure location:** The development and testing infrastructures are located at RAL-STFC, Didcot, UK
- **Team composition:** two developers (54 PMs over 4 years for both GOCDB and Accounting Repository) are responsible for the development of the tool
- **Distribution of work inside the team:** internal task redistribution is planned in the near future but currently waiting on staff to finish their current project
- **Code repository used:** SVN is used to manage the code
- **Build system:** KOJI [13] is used as build system
- **Bug/task tracking facilities:** a Savannah project hosted at CERN [12]
- **Internal communication channel:** the savannah project and the gocdb-admins<at>mailtalk.ac.uk mailing list
- **Support communication channels:** a dedicated GGUS support unit and the gocdb-admins<at>mailtalk.ac.uk mailing list
- **Testing and validation procedures:** Development is done on a dedicated machine. Functional and compatibility tests are done on there. The code is then packaged and installed on a clean test machine where all tests are repeated. If all tests are successful the release is ready for production. If not the code goes back to development. For the GOCDB programmatic interface, a set of test clients has been implemented that connect to the PI and retrieve results. These are used to verify the functionalities of the PI
- **Distribution facilities:** code is distributed through SVN

2.1.4 Accounting Repository PT and development infrastructure

The EGI accounting repository stores information about the usage sites resources within the EGI production infrastructure. It receives records from accounting probes running in the sites and exports accounting information to consumers (i.e. the accounting portal for visualization).

Currently the accounting repository runs as a central instance only, but development are foreseen to have a distributable version that can be installed at a regional or NGI level.

Details on the architecture and implementation are available at [14] and [15].

Description of the production and testing infrastructure:

- **Infrastructure location:** same team developing the GOCDB at STFC (54 PMs over 4 years for both components)
- **Team composition:** same team developing the GOCDB at STFC

- **Distribution of work inside the team:** internal task redistribution is planned in the near future but currently waiting on staff to finish their current project
- **Code repository used:** SVN is used to manage the code
- **Build system:** KOJI [13] is used as build system
- **Bug/task tracking facilities:** a savannah project hosted at CERN [32]
- **Internal communication channel:** the APEL-SUPPORT<at>JISCMail.AC.UK mailing list
- **Support communication channels:** GGUS helpdesk dedicated support unit and the APEL-SUPPORT<at>JISCMail.AC.UK mailing list
- **Testing and validation procedures:** use a testing infrastructure to simulate the production environment. The same checks that are done to monitor the production service are used to test the development one
- **Distribution facilities:** not distributed for now

2.1.4.1 Accounting for different resource types

The work needed for task TJRA1.4 (accounting for different resource types, see section 3) will be spread over four partners: LUH (18 PMs), INFN (26 PMs), SFTC (27 PMs), CESSGA (18PMs).

Details on the development infrastructure are still not available since the task will start on the second year of the project. The project milestones that will define the workplan for this task will be the MS706 (due month 13), MS709 (due month 25), MS711 (due month 37).

2.1.5 Accounting Portal PT and development infrastructure

The Accounting Portal is a graphical frontend for the production infrastructure accounting data. Accounting statistics are available for different kinds of users: i.e. end users, VO managers, site administrators.

The central accounting portal will be extended to support regional deployments of NGIs accounting systems and national deployment of the accounting repository. The Accounting Portal will be distributed to NGIs that wish to install it, but also new regional views on the central instance will be created for NGIs not running their own instance of the portal.

Implementation and architecture details can be found at [16].

Description of the production and testing infrastructure:

- **Infrastructure location:** The developers and development infrastructure are located at CESSGA
- **Team composition:** 4 people – 23 PMs
- **Distribution of work inside the team:** 3 PMs for the coordination, 20 PMs for the development
- **Code repository used:** SVN [18] is used to manage the code
- **Build system:** not built
- **Bug/task tracking facilities:** Internal CESSGA RT [17] helpdesk is used as bug/task tracking system
- **Internal communication channel:** Internal CESSGA RT
- **Support communication channels:** dedicated GGUS support unit and the grid-admin<at>cerga.es mailing list
- **Testing and validation procedures:** Development version is available for preview and testing on a dedicated instance. The testing procedures are in general composed by two step:



- tests performed by the development group
- final review performed by the requestor of the new functionality
- **Distribution facilities:** not distributed for now

2.1.6 Service Availability Monitoring PT and development infrastructure

The Service Availability Monitor (SAM) is the system that will be used to monitor the resources within the production infrastructure. It includes the following components:

- The test framework: based on the NAGIOS system [19] 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 DataBase (MRDB)
- Message bus to publish the monitoring results
- A visualization tool: MyEGI

The SAM system is a distributable package and its regionalization was almost completed during EGEEIII [20]. The validation process of the regional instances is still ongoing and progress is tracked on this wiki page [21].

Description of the production and testing infrastructure:

- **Infrastructure location:** Distributed across two partners: SRCE and CERN
- **Team composition:** : SRCE, two people for a total of 15 PMs, CERN two people for a total of 18 PMs
- **Distribution of work inside the team:** : SRCE is responsible for the NCG component, CERN is responsible for the DB components and MyEGI, the message bus configuration is done in conjunction with a team located at AUTH (see section 4.1.7)
- **Code repository used:** code is accessible through an SVN repository [27]
- **Build system:** a KOJI installation [13] [24], provided by AUTH, is used as build system.
- **Bug/task tracking facilities:** JIRA [22] [23], hosted at CERN, is used as bug/task tracking system
- **Internal communication channel:** JIRA [22] [23], hosted at CERN
- **Support communication channels:** dedicated GGUS support unit and the regional-nagios-admins[at]cern.ch mailing list (will be ported to an egi.eu domain)
- **Testing and validation procedures:** Each release is bundled in a JIRA version which has a single task tracking the validation of that release through a dedicated testbed where testing is done to check that the bugs are fixed.
- **Distribution facilities:** YUM repositories are created by the KOJI building system at AUTH [25] and synchronized off-site to Manchester [26] - both sites can act as distribution facilities

2.1.7 Message Bus Configuration PT and development infrastructure

Within the activity a group located at AUTH will work to ease the integration of the tools with the broker network that EGI is running. This broker network is based on ActiveMQ/FUSE [28] and even if its development or maintenance is not responsibility of the JRA1 activity, work will be carried on to provide configuration and tools that will maintain a message bus infrastructure that fit each application's requirements. Some tools, as SAM for example, are already integrated with the ActiveMQ broker network and their message bus configuration will be kept updated as the EGI production broker will evolve.



Description of the production and testing infrastructure:

- **Infrastructure location:** AUTH
- **Team composition:** three people for a total of 12PMs
- **Distribution of work inside the team:** Development related communication are handled within JIRA
- **Code repository used:** code is accessible through the Manchester's SVN repository [27]
- **Build system:** KOJI (at AUTH) [25] as the building system
- **Bug/task tracking facilities:** JIRA (at CERN) [22] [29] for bug/task tracking system
- **Internal communication channel:** JIRA
- **Support communication channels:** Production message bus network maintenance and operation communication are handled at the following mailing list hosted at AUTH: prod-broker-network<at>grid.auth.gr
- **Testing and validation procedures:** follows the same testing procedure done for the SAM system (section 4.1.6)
- **Distribution facilities:** YUM repositories at AUTH and Manchester [25][26]

2.1.8 Metrics portal PT and development infrastructure

The metrics portal automatically collects a set of metrics that can help to measure project performance and keep track of its evolution displaying all of them in a single place. The portal collects information from different sources like GOCDB, GGUS, GridView, etc. using various connectors and stores it into a local database. Data are then available through web pages that can be exploited to create statistical reports in few seconds. No regional model is foreseen, the portal will run on a central installation only.

Architecture and implementation details can be found at [30].

Description of the production and testing infrastructure:

- **Infrastructure location:** The development group and the testbeds are located at CESGA
- **Team composition:** 3 people – 10 PMs
- **Distribution of work inside the team:** 2PM for the coordination and 8PM for the development
- **Code repository used:** SVN for managing the code
- **Build system:** not built
- **Bug/task tracking facilities:** Internal CESGA RT is used as bug/task tracking system
- **Internal communication channel:** grid-admin<at>cesga.es mailing list
- **Support communication channels:** grid-admin<at>cesga.es mailing list
- **Testing and validation procedures:** include internal review done by the developers of the new functionalities and a preview instance made available to the functionality requestors
- **Distribution facilities:** not distributed

	Ops Portal	GGUS	GOCDB	Acc. Repos.	Acc. Portal	SAM	Message Bus	Metrics Portal
Location	CNRS, France	KIT, Germany	RAL-STFC, UK	RAL- STFC, UK	CESGA,S pain	CERN, Switzerland	AUTH, Greece	CESGA, Spain

						SRCE, Croatia AUTH, Greece		
Code repository	SVN	SVN	SVN	SVN	SVN	SVN	SVN	SVN
Build system	Maven Symfony scripts	Not built	KOJI	KOJI	Not built	KOJI	KOJI	Not built
Bug/Task tracking	Forge	savannah	savannah	savannah	RT	JIRA	JIRA	RT
Distrib. facilities	SVN	Not distrib.	SVN	Not distrib. for now	Not distrib. for now	YUM	YUM	Not distrib.
Internal comm. channels	Mailing list	Mailing list	Mailing list/Savannah	Mailing list	RT	JIRA	JIRA	Mailing list
Support comm. channels	Webform GGUS	Mailing list GGUS	Mailing list GGUS	Mailing list GGUS	Mailing list GGUS	Mailing list GGUS	Mailing list	Mailing list

Table 1 - Summary of the PTs development infrastructure



3 INTERACTIONS OF THE PRODUCT TEAMS

OTPTs need to interact with each other inside the activity, as well as with other project activities, primarily with SA1, the main customer of the operational tools.

Interactions among OTPTs is guaranteed by the activity management through periodic phone conferences and face to face meetings. Activity progress will be tracked using EGI.eu facilities such as the RT system [31], and will be reported in periodic reports. A dedicated mailing list (under the egi.eu domain) is available for activity as well.

Interaction with external bodies is fundamental in order to get input, feedback and new requirements for the developed tools. Two advisory groups are foreseen by the project description of work: the Operational Tools Advisory Group (OTAG) and the User Services Advisory Group (USAG):

- Being composed by representatives from the operation community, from the middleware developers and from the JRA1 activity, the OTAG will be the main supervisory group for the development progress and the place where technical discussion about the evolution of the tools will take place
- The USAG has representatives from the EGI user communities and will focus on the requirements for the complete set of services run by the project but could also impact on the operational tools - for example end user requirements could be addressed to the GGUS Helpdesk

In order to create the proper schedule for the development it will be important that the outcome of the advisory groups will be a single prioritized table. The prioritization is an important step and possible conflicts should be resolved when multiple requirements from different groups impact the same tool. This is not expected to happen frequently and will be analyzed case by case by the management of the activity together with the advisory groups, possibly escalating the problems to the Activity Management Board if needed.

The development progress of new features, requested and approved by the advisory groups, will be tracked using the project tracking system [31].

The representation of the activity within other project bodies and the reporting on the status of the activity is a responsibility of the activity manager.

4 OUTPUT OF THE PRODUCT TEAMS

Each tool will be released as a standalone package. PTs are autonomous in the development, but the release schedule and roadmap will be discussed and agreed both internally and with other project actors if needed (i.e. SA1).

Testing and documenting the released packages are responsibilities of the PTs under the supervision of the activity management. If a new release for an operational tool affects other tools or middleware installed in the production infrastructure, then a test plan will be discussed among the PTs and/or with the SA1 activity. To ease the information exchange and to discuss test plans a release procedure, which is described below and depicted in Figure 2, will be applied by every PT.

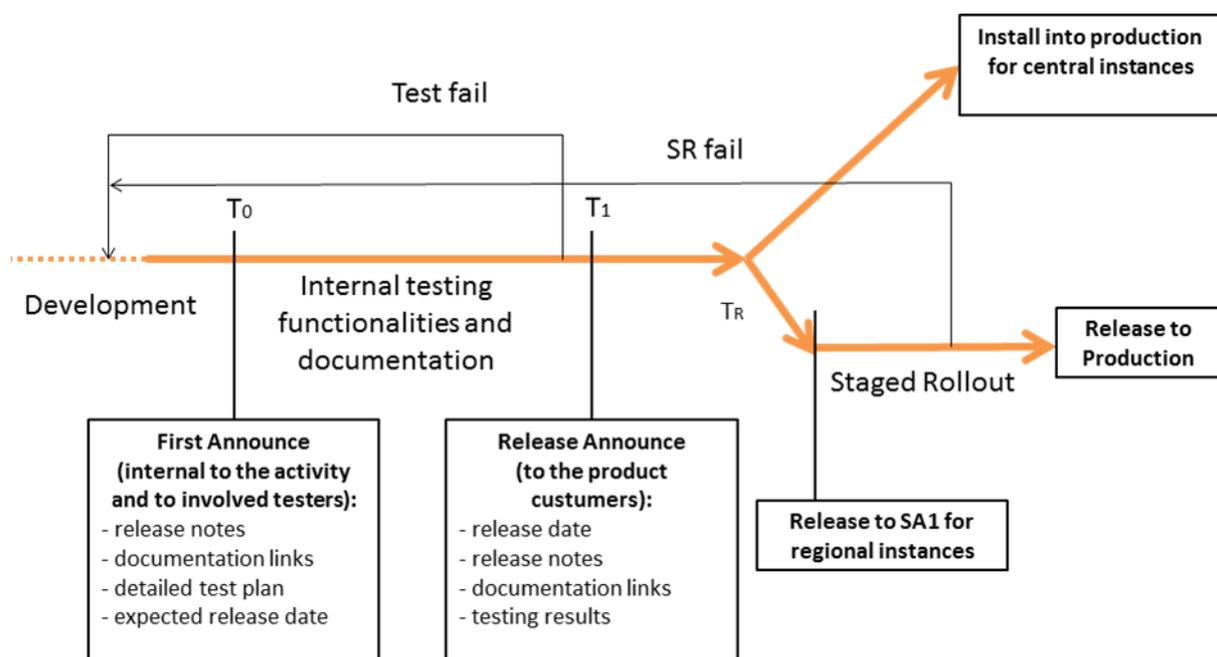


Figure 2 - Operational Tools Release Procedure

When the development of a new release is completed (T_0) a first announce is broadcasted to all the PTs and to all the actors that should be involved in the release testing. The announce has to contain release notes, documentation links, a detailed test plan, an indication of the expected release date and all the information needed by the conformance criteria set by the SA2 activity for the software providers of the project. More information on conformance criteria can be found on the project milestone MS503. The expected release date and the kind of testing will depend on each specific release and on its importance.

After the T_0 announce the testing phase will take place according to the test plan. If a test fails a report will be produced and the release sent back to development to restart the cycle. Tests will



include a documentation review and a documentation update if needed. The test phase can be performed internally to the PT if no other tools or services are affected. When all the tests are passed and no more development/testing cycles are needed the testing phase is concluded (T1) and a second release announce will be broadcasted to the consumers of the new release (i.e. SA1) using the communication facilities offered by the Operation Portal. This second announce will contain the actual release date (T_R), the release notes, the documentation links and a document describing the testing phase details. The release can result in an immediate installation on the production instances for the centralized tools or will follow a deployment process, with a possible initial testing phase on a selected number of production instances according to the SA1 needs (StagedRollout, for further details refer to the project milestone MS402).

Each tool will maintain its own documentation on the web and links to these web pages will be maintained on the project wiki by the activity management in order to have a single access point to all the tools documentation. The activity management will also supervise the need and the editing of cross tools documentation to support the integrated usage of multiple components.

4.1 Strengthening the test phase and monitoring the software quality

Testing of the released software is a responsibility of the product teams, however the manpower available for each OTPT does not allow for fully independent testing. In almost all PTs the developers act also as testers, even if, when possible, not the same person performs both activities on the same piece of code. To mitigate this situation it will be important for a PT to discuss the testing plan within the whole jra1 activity to agree contribution to testing from other PTs. Moreover in case of major or particularly important releases contribution can also be found outside the activity, mainly inside the operation community.

The quality of the operational tools releases will be constantly monitored through a set of metrics that will include, in example, the number of bugs found in certification and in production, as well as the time needed to fix critical and non critical bugs by each PT. This metrics are currently under definition. If the monitoring activity will show that the software quality of some operational tools needs to be improved it will reported to the projects management and actions will be undertaken in order to strengthen the testing phase, i.e. increase the number of external and independent testers to be found both internally to the partners developing the tools and externally within the community.

Given the limited manpower the prioritization work performed by the supervising bodies will be particularly important in order not to waste developing and testing efforts.



5 FUTURE OF THE PRODUCT TEAMS

It results from section 4 that each product team has its own development infrastructure and organization with a central coordination provided by the JRA1 activity management. The development infrastructures are those inherited from the previous EGEE projects and it is not foreseen, in the near future, their restructuring in order to have a standardized development model. Given the present number of PTs, this model of autonomous teams with a central coordination seems to be sustainable. Moreover the decision of not modifying the existing development infrastructures resulted in no time spent for organizational matters at the beginning of the project, allowing each team to be immediately productive.

New product teams can be created if needed. In particular the task JRA1.4, accounting for different resource types, to be started on the second year of the project, could demand the creation of a new product team. Such process will be agreed with the project management and within the activity in order to harmonize the new PT with the existing ones. If in the long term the increased number of autonomous PTs will result in a not sustainable development infrastructure the decision about the standardization of the infrastructures can be reviewed.



6 CONCLUSIONS

In this milestone the development and testing infrastructure for the operational tools under the responsibility of the JRA1 activity was described. The infrastructure is distributed across the partners participating to the activity and is composed by different (in terms of hardware resources, bug/task tracking systems, communication channels and deployment facilities etc.) sub-infrastructures. All the OTPTs already started their work, even if the staff definition is still ongoing at RAL (waiting for people to switch from their previous activity) and at CESGA (hiring process to be completed). Each product team is autonomous in the development but the requirements workflow and the release process used during the whole project lifetime will be common to all the tools. The activity management will supervise the development roadmap and will provide the needed coordination effort in order to handle the distributed nature of the development.

7 REFERENCES

R 1	Symphony open source PHP framework: http://www.symfony-project.org/
R 2	Operation Portal homepage and general documentation page: https://operations-portal.in2p3.fr https://cic.gridops.org/index.php?section=roc&page=generaldoc
R 3	Apache Maven build manager for Java projects: http://maven.apache.org/
R 4	Operation Portal Forge project at IN2P3: https://forge.in2p3.fr/projects/show/opsportaluser
R 5	Operation Portal CVS at IN2P3: https://cvs.in2p3.fr/operations-portal/package/
R 6	BMC Remedy IT service management Suite: http://www.remedy.com/
R 7	EGI Helpdesk home page and main documentation page: http://ggus.org https://gus.fzk.de/pages/docu.php
R 8	EGI Helpdesk shopping list hosted by CERN Savannah: https://savannah.cern.ch/projects/esc
R 9	EGI Helpdesk ongoing wish list: https://gus.fzk.de/pages/owl.php
R 10	GOCDDB architecture details: http://goc.grid.sinica.edu.tw/gocwiki/GOCDDB4_Architecture
R 11	GOCDDB technical documentation: http://goc.grid.sinica.edu.tw/gocwiki/GOCDDB_Technical_Documentation
R 12	GOCDDB project hosted by CERN Savannah: https://savannah.cern.ch/projects/gocdb/
R 13	KOJI build system homepage: https://fedorahosted.org/koji/
R 14	Accounting Repository documentation main page: http://goc.grid.sinica.edu.tw/gocwiki/ApelHome
R 15	The Accounting Infrastructure in the EGEE projects: http://www3.egee.cesga.es/gridsite/accounting/CESGA/links/acct_ibergrid06_final14.pdf
R 16	Accounting Portal architecture and implementation details: https://edms.cern.ch/document/726137
R 17	Request Tracker homepage: http://bestpractical.com/rt/

R 18	Apache Subversion homepage: http://subversion.apache.org/
R 19	NAGIOS alarm system homepage: http://www.nagios.org/
R 20	Production service status at the end of the EGEEIII project: https://edms.cern.ch/document/1060571
R 21	Wiki page at CERN tracking the status of the SAM regionalization: https://twiki.cern.ch/twiki/bin/view/EGEE/ExternalROCNagios
R 22	JIRA bug & issue tracker home page: http://www.atlassian.com/software/jira/
R 23	SAM project hosted by the JIRA instance at CERN: https://tomtools.cern.ch/jira/browse/SAM
R 24	KOJI instance at ATH: https://koji.afroditi.hellasgrid.gr/koji/
R 25	RPMs distribution facility at AUTH: http://rpm.hellasgrid.gr/mash/
R 26	RPMs Distribution facility at Manchester: http://www.sysadmin.hep.ac.uk/rpms/egee-SA1/
R 27	SAM SVN instance at Manchester: http://www.sysadmin.hep.ac.uk/svn/grid-monitoring/
R 28	ActiveMQ message system homepage: http://activemq.apache.org/
R 29	Messaging project hosted by the JIRA instance at CERN: https://tomtools.cern.ch/jira/browse/MIG
R 30	Metrics Portal architecture and implementation: http://www3.egee.cesga.es/gridsite/accounting/CESGA/links/paper_metrics_iber2010.pdf
R 31	RT facility hosted by the egi.eu portal: https://rt.egi.eu/rt/index.html
R 32	Accounting Repository project hosted by CERN Savannah: https://savannah.cern.ch/projects/apel
R33	GOCD home page: https://goc.gridops.org/
R34	Accounting Portal home page: http://www3.egee.cesga.es/gridsite/accounting/CESGA/egee_view.php
R35	Metrics Portal homepage: http://www3.egee.cesga.es/gridsite/metrics/CESGA/egee_view.php