EGI-InSPIRE

EGI OPERATIONS ARCHITECTURE
INFRASTRUCTURE PLATFORM AND COLLABORATION PLATFORM INTEGRATION

EGI-InSPIRE DELIVERABLE: D4.6

Document identifier:	EGI-docid-V9.doc
Date:	**01/10/2012**
Activity:	**SA1**
Lead Partner:	**EGI.eu**
Document Status:	**DRAFT**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/1309

Abstract

This document presents how services that are part of the EGI operations portfolio, which contribute to the implementation of the EGI Infrastructure Platform and Collaboration Platform, can be adopted and/or extended by external Research Infrastructures (RIs) to support integrated operations. The adoption of EGI operations tools, deployed and developed within EGI-InSPIRE, allows the reuse of existing solutions that address the typical needs of any distributed RI.

I. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Tiziana Ferrari	EGI.eu/SA1	
Reviewed by	**Moderator:** **Reviewers:**		
Approved by	**AMB & PMB**		

II. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	01-10-2012	V1.0	S. Burke, T. Ferrari, P. Solagna/EGI.eu
2			
3			

III. 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.

IV. 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

V. TERMINOLOGY

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

The following list provides a set of terms that are used in this document.

- **EGI Platform model.** The EGI Platform model refers to business models that may emerge by utilising any of the IT platforms that are described in the EGI Platform architecture.

- **EGI Platform architecture.** Describes how the individual platforms (see below) are embedded in the EGI ecosystem, and how they are technically integrated with the current EGI production infrastructure.

- **EGI Infrastructure Platform.** The EGI Infrastructure Platform comprises of IT Infrastructure and IT Services that are required by all Research Communities that are part of the EGI ecosystem in order to deliver community-specific services and infrastructure.

- **EGI Collaboration Platform.** The EGI Collaboration Platform provides IT Infrastructure and Services that facilitate collaboration between Research Communities without being a core infrastructure service for Research Communities.

- **EGI Community Platform.** EGI Community Platforms (there may be more than one) consist of services that are specific to the respective community's needs.

VI. PROJECT SUMMARY

To support science and innovation, a lasting operational model for e-Science is needed − both for coordinating the infrastructure and for delivering integrated services that cross national borders.

The EGI-InSPIRE project will support the transition from a project-based system to a sustainable pan-European e-Infrastructure, by supporting 'grids' of high-performance computing (HPC) and high-throughput computing (HTC) resources. EGI-InSPIRE will also be ideally placed to integrate new Distributed Computing Infrastructures (DCIs) such as clouds, supercomputing networks and desktop grids, to benefit user communities within the European Research Area.

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support will also be given to the current heavy users of the infrastructure, such as high energy physics, computational chemistry and life sciences, as they move their critical services and tools from a centralised support model to one driven by their own individual communities.

The objectives of the project are:

1. The continued operation and expansion of today's production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.

2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.

3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.

4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.

5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.

6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The EGI community is a federation of independent national and community resource providers, whose resources support specific research communities and international collaborators both within Europe and worldwide. EGI.eu, coordinator of EGI-InSPIRE, brings together partner institutions established within the community to provide a set of essential human and technical services that enable secure integrated access to distributed resources on behalf of the community.

The production infrastructure supports Virtual Research Communities (VRCs) − structured international user communities − that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.

VII. EXECUTIVE SUMMARY

The foundation of the EGI Platform architecture is the definition of the term *platform*: "a composition of IT Infrastructure and IT Services that together enable independent solution providers to build other technologies or processes, or both, on top of it". This document describes how several operations services that contribute to the implementation of the *Infrastructure Platform* and the *Collaboration Platform* of EGI [M506], can be used by external RIs as building blocks to implement an own operations service portfolio. Different configurations are possible depending on the level of integration with EGI services that needs to be implemented. The document also provides information about the suitability and the extensions needed in order to support Federated Cloud resource provisioning.

The EGI service registry – GOCDB – is software independent and can be deployed by any RI (both cloud- and/or grid-based) to register service end-points, contact information, scheduled maintenance etc. Information from about multiple independent RIs can be hosted centrally by a single GOCDB instance, thanks to the capability of associating service end points to different target infrastructures. Alternatively, a stand-alone dedicated GOCDB instance can be deployed.

The Accounting Infrastructure of EGI relies on three major components: service-level probes and publishers to extract usage information from log files, a broker network to exchange usage records, and a central database to collect information and present views through an accounting portal. In case of user communities coupling resources from multiple RIs, a unified view of accounting information is often needed. This requires the collection of accounting data from multiple RIs. For a fully interoperable scenario, records need to conform to the structure recommended by the applicable standards. The exchange of records between authorized publishers through the message bus allows multiple RI consumers to gather information and store it locally. Accounting of cloud resource usage requires the development of probes extracting the relevant metrics locally; the EGI central accounting repository based on APEL is being extended in order to collect usage information for many different resource types.

The Service Availability Monitoring (SAM) is a framework for the analysis of the status of public service interfaces. It can be easily extended to probe heterogeneous services according to the RI needs. SAM together with GOCDB, which provides topology information, and the Operations Dashboard for the display of alarms, provide a complete solution to the problem of infrastructure monitoring and visualization of failures at three different aggregation levels: Resource Centre, RI and federation of RIs.

The EGI service discovery service, based on GLUE and LDAP, supports service discovery, service selection, monitoring, oversight and diagnostic. The service discovery service requires the availability of an abstract resource representation model. GLUE 2 is being extended to support publication of cloud resources.

Finally, the user support infrastructure of EGI is inherently distributed consisting of various topical and regional helpdesk systems that are linked together through a central integration platform, the GGUS helpdesk. This central helpdesk enables formalized communication between the submitter of the incident record and all partners involved in user support by providing an interface to which all other tools can connect and enabling central tracking of a problem, independent of the origin of the problem and the tool in which the work on the problem is done. The interlinking of all ticket systems already in place within EGI enables to pass incident records from one system to the other in a way that is transparent to the user, and makes the integration with other RI helpdesk systems easy. EGI can offer to user communities and RIs a stand-alone dedicated helpdesk as software as a service (it is based on the xGUS technology), which is fully integrated with external service desks according to

their needs, The integration of GGUS with the external service desks operated by other RIs allows the provisioning of a seamless support structure, where incident records can be migrated to one helpdesk to the other depending on the RI concerned by the problem and the type of support needed.

The adoption of EGI operations tools, deployed and developed within EGI-InSPIRE, allows the reuse of existing solutions that address the typical needs of any distributed RI. These needs are typically resource and technology independent. EGI can share these building blocks in a variety of configuration options.

EGI can operate these tools in a software-as-a-service fashion for other RIs, can service RIs through the existing production instances of EGI, or alternatively provide the tools as open-source software for the local deployment of independent instances that can be easily integrated with existing EGI production services.

TABLE OF CONTENTS

1 INTRODUCTION

The foundation of the EGI Platform architecture is the definition of the term *platform*: "a composition of IT Infrastructure and IT Services that together enable independent solution providers to build other technologies or processes, or both, on top of it" [M506].

The EGI Platform comprises three distinct types of platforms that serve different purposes. The *EGI Infrastructure Platform* enables consistent access to a large federated distributed computing infrastructure such as virtualised compute, storage and network resources, and supplemental services such as Information Discovery, Accounting, Monitoring and Notification. It enables platform integrators to utilise this solid base to build different higher-level infrastructures (virtual research environments), for example targeting the requirements emerging from a wide variety of research infrastructures as documented by the ESFRI. The Infrastructure Platform interfaces can also be extended or integrated with equivalent ones operated by external infrastructure providers, so that different RIs can be seamlessly operated.

Extending the EGI Infrastructure Platform, the *EGI Collaboration Platform* facilitates synergies between Research Communities by providing services that are common across user communities without being domain-specific or critical to the operation of EGI itself. Services such as federated AAI (for Platform Users), Service Desks or meeting planning systems are examples of such facilitating collaborative services.

On top of the EGI Infrastructure Platform, while making use of the EGI Collaboration Platform where required, any numbers of *EGI Community Platforms* provide domain specific access to the distributed EGI computing infrastructure and integrating elements of its platforms into domain specific Virtual Research Environments.

This document presents how services that are part of the EGI operations portfolio [D4.3] that contribute to the implementation of the EGI Infrastructure Platform and Collaboration Platform, can be adopted and/or extended by external Research Infrastructures (RIs) to support integrated operations.

Services of the Infrastructure Platform are presented in Chapter 2, which provides information about how integrated RIs can share topology information in GOCDB (Section 2.1), exchange accounting information (Section 2.2), be seamlessly monitored (section 2.3 and 2.4) and discovered (section 2.5). Chapter 3 describes the potential role of the EGI helpdesk system (GGUS) and xGUS as collaboration service facilitating support to multiple user communities, making use of different RIs and relying on the deployment of different software technologies. Chapter 3 describes how various support configuration can be implemented through different helpdesk integration strategies. Finally, Chapter 4 concludes the document.

2 INFRASTRUCTURE PLATFORM

This chapter provides information on how external RIs can deploy several components of the EGI Collaboration Platform that any RI can use as building blocks to implement an operations service portfolio, namely: the GOCDB service registry, accounting, Service Availability Monitoring, the Operations Dashboard and the Information discovery service.

Various integration scenarios are illustrated that allow a RI to implement a set of operational services that are either fully integrated with the EGI ones, or autonomously operated as stand-alone services.

2.1 GOCDB

Grid Configuration Database[1] (GOCDB) [GOCDB] contains general static information about the Resource Centres (also known as "sites") participating to one or more infrastructures, and most importantly, it provides topology information about them by defining the respective list of service end-points. GOCDB allows sites to store, to maintain and view the topology of the production infrastructure and the basic information about the respective resources within it. GOCDB can provide information for multiple target infrastructures and is independent from the services being registered. Examples of information provided are:

- Participating Resource Providers (National Grid Initiatives, European Intergovernmental Organizations), the respective Operations Centres and the related information (countries, contact information etc.).

- Resource Centres contributing resources to the infrastructure including management, technical and security related contact points.

- Resources and services, including scheduled intervention plans and service status information (e.g. certification, production and monitoring status).

- Participating people and their roles within EGI operations, where roles limit operations that particular group of people can perform.

Any type of service can be registered (Community Platform services, Infrastructure Platform services, etc.). Service end-points are described with the following information:

- Service Type**:** a unique name that identifies the type of software component deployed on a Grid.

- Service Endpoint**:** is a deployed instance of a named service type

- Endpoint Location: a Service Endpoint may optionally define an Endpoint Location which locates the service (URL).

Besides providing a central management tool to view and define production state, downtimes and maintenance status and whether a resource needs monitoring, it shall in essence depict what services are running where and who to contact for certain type of issues.

---

[1] http://goc.egi.eu/

A management interface provides information about a resource through the certification process. The history and details of the certification status transitions and other state transitions like site decertification and suspension are desirable additional information.

GOCDB is also referent database for all other operational tools, providing all the relevant data about Resource Providers, Resource Centres, services and administrators responsible for resources and services.

All e-Infrastructures have to deal with topology and operational information like resource distribution, contact lists and downtime declarations. Storing, maintaining and publishing this information properly is one of the key elements to successful grid operations. GOCDB as a central repository that hosts this information and makes it available to users and client tools is an answer for this need. This repository is for each RI an authoritative primary source of information for operations, monitoring, accounting and reporting. GOCDB is a central database, but it is designed in a way that allows for including resources that are part of other infrastructures. This can be done by operating independent instances of the tool, through the usage of scoping, i.e. by declaring resources not being part of a RI, or by associating sites and end-points to different target infrastructures.

Various integration scenarios are possible, depending on the needs. These are described in the following sections.

2.1.1 Standalone GOCDB instances

Standalone GOCDB independent instances can be operated, each GOCDB instance serving a different RI. In such case, the GOCDB instance is fully dedicated to a given RI. This integration scenario is suitable for disjoint infrastructures that comprise different sites and service end-points, otherwise operating multiple service instances will require information about shared service to be duplicated in different GOCDB instances.

Independent GOCDB instances can be autonomously operated directly by infrastructure providers, or provided by EGI.eu on behalf of them.
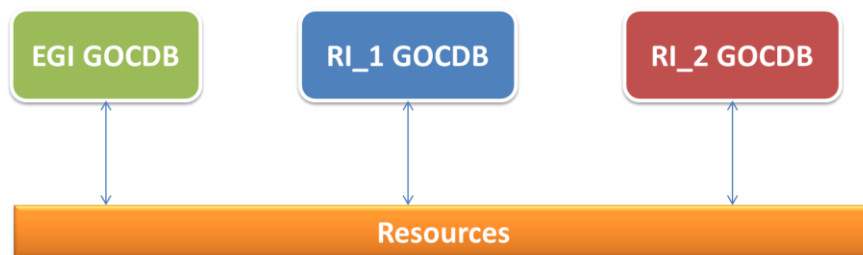
Figure 1. Different RIs can rely on dedicated GOCDB instances. In this integration scenario, information about service end-points and/or Resource Centres belonging to different RIs at the same time has to be replicated in multiple GOCDB instances.

2.1.2 Shared GOCDB hosting different target infrastructures

A single shared GOCDB instance can service multiple RIs by associating sites (and the corresponding service end-points) to different target infrastructures. This is applicable when one site and its service end-points belong the same RI (Figure 2).

On the other hand, if a single site hosts service end-points belonging to different RIs, groupings can be used (Figure 3). A service group is an arbitrary grouping of existing service endpoints that can be distributed across different physical sites. Currently, any GOCDB user can create their own service groups (everything is logged, including who created the service group).
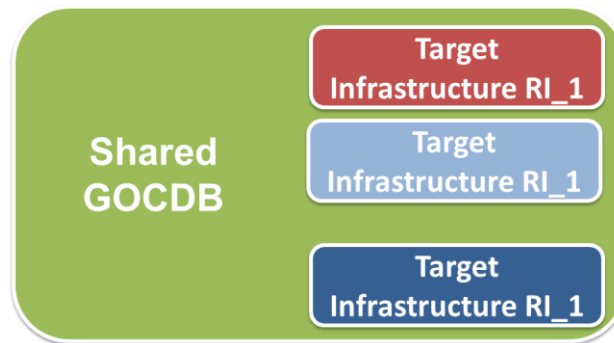
Figure 2. A singe GOCDB instance can host sites of different target infrastructures.
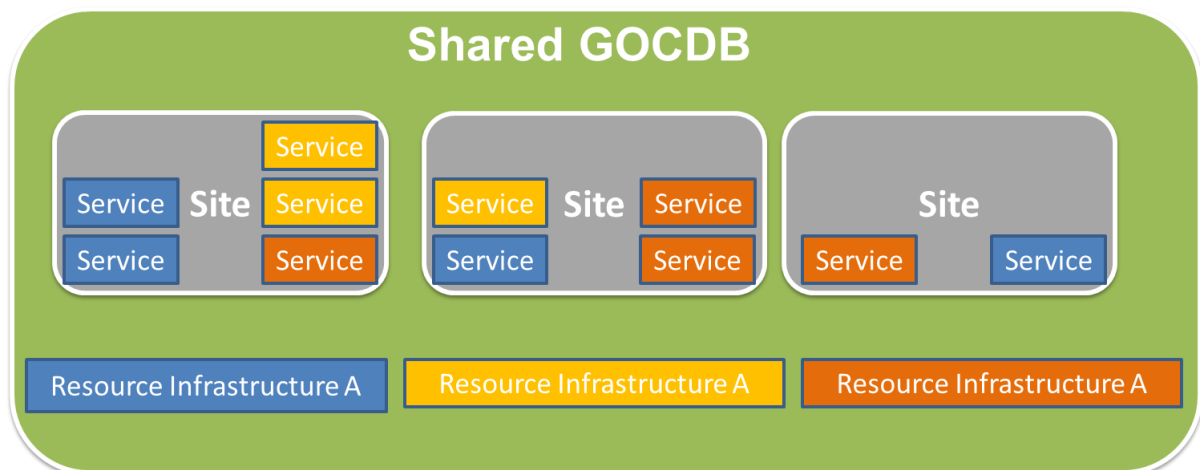
Figure 3. Support of different RIs through a single shared GOCDB instance and the usage of service end-point grouping.

2.1.3 Community Platform services

Community Platform services can be specific to a given user community. In order to record these, the CUSTOM service type can be used. This feature allows for including software whose scope may be limited to a specific organizational unit (e.g. RI or VO), or software which has been customised so that its functionality is no longer standard. The service type name reflects the usable scope of deployment. All data are stored in one database and provided through a common web interface.

2.1.4 Federated Cloud support

GOCDB can be easily adopted to allow the registration of cloud services, and to keep track of information for the virtual resources, to be used by users or other operational tools such as SAM and accounting.

2.2 Accounting

The usage of the EGI computing resources by the users is accounted at service level and the accounting data is stored in a central repository, in order to calculate the overall usage statistics of the distributed resource infrastructure, aggregating the contribution from the different resource centres.

The accounting infrastructure is composed by three main components:

- **Service-level probes and publishers**: component that parses the logs of the grid computing element and the batch system, extracting the accounting information and creating the usage records to be stored centrally

- **Brokers network**: messaging transport infrastructure that takes care of the delivery of the usage records from the publisher to the repository

- **Accounting repository and portal**: central database to contain the usage records produced by the production sites part of the infrastructure. An Accounting Portal is available to access the information stored in the accounting repository, includes different views for the sites, users and VO Managers.

The EGI Accounting Infrastructure currently collects CPU accounting records from sites and/or grid infrastructures and summarizes the data by site, date (especially by month), VO, and user. This summary data can be displayed in a central Accounting Portal[2] by dynamic queries on the parameters above at any level of the hierarchical tree structure which defines EGI and its partner grids. The core EGI Accounting Infrastructure is based on APEL [APEL].

Other accounting systems have to interface to APEL to publish the data in the central repository. The Secure Stomp Messenger (SSM) [SSM] is the messaging system used by APEL to securely transmit messages. It is written in Python and uses the STOMP protocol. For a unified access to VO accounting information, RIs can publish VO usage records centrally into the EGI accounting repository, or into a different repository of choice (see Figure 4).

The bulk of existing Resource Centres collect data from their batch systems (e.g. LSF, Torque, GE), which are joined with information about the job's user grid credentials and published to the central APEL repository. Other partner grids (currently Open Science Grid, and CERN), already publish Usage Records through SSM, while to date others are migrating to the new transport protocol.

The OGF UR Working Group (UR-WG) is revising the standard for compute accounting [CAR], and is considering a proposal from EMI for a UR for storage accounting [STAR]. It is anticipated that this will be integrated into the same APEL infrastructure once implemented on the relevant storage products. EMI also has a group reviewing the implementations of the OGF UR for compute accounting to agree on the semantics of the existing UR and existing common extensions and possibly propose further extensions.

---

[2] http://accounting.egi.eu/egi.php

2.2.1 Accounting integration

When looking at the accounting interface as the interface between the accounting services of different interoperating infrastructures the main aim is to enable all the accounting data of a VO to be collected in one place for a unified view. While participant RIs currently publish all of their VOs data, partner grids can publish information for a subset of VOs (e.g. OSG) according to the VO requirements.

Other RIs who wish to publish accounting data need to:

- Adopt an OGF standard Usage Record scheme and an EGI-compatible profile. A draft exist of the EGI Accounting Profile [EAP] for the new emerging OGF standards (Compute Accounting Record [CAR] and STorage Accounting Record [STAR]), which will be finalized by the end of 2012.

- Define a structure for their service end-points in GOCDB (or equivalent) that can be used by the accounting portal to display the data. The minimum requirement is a flat set of site names, used in the accounting records (e.g. for OSG these data are obtained from MyOSG[3]). This is required only if the EGI Accounting Portal is the tool of choice for visualization of the data.

- Extract data from their accounting system and group data by site/VO/User/FQAN/ month and create each group into a 'summary record' meeting the APEL definition. Experience shows that for accounting systems using the OGF-UR this is a simple transformation.

- Register the publisher with APEL as the APEL Repository only accepts accounting records from registered Resource Centres.

- Publish the records into EGI's ActiveMQ[4] Message Bus using the agreed encryption framework and SSM. ActiveMQ is an open source (Apache 2.0 licensed) message broker which fully implements the Java Message Service 1.1 (JMS). The APEL repository will accept the records into a holding container from where they will be merged with the summaries from other RIs and the summary produced by APEL from the job records it has received. Currently, the master summary is rebuilt from scratch several times per day. Each time it uses the last set of summaries received from each grid.

- From the master summary table, the data are then exported to the Accounting Portal where they can be viewed.

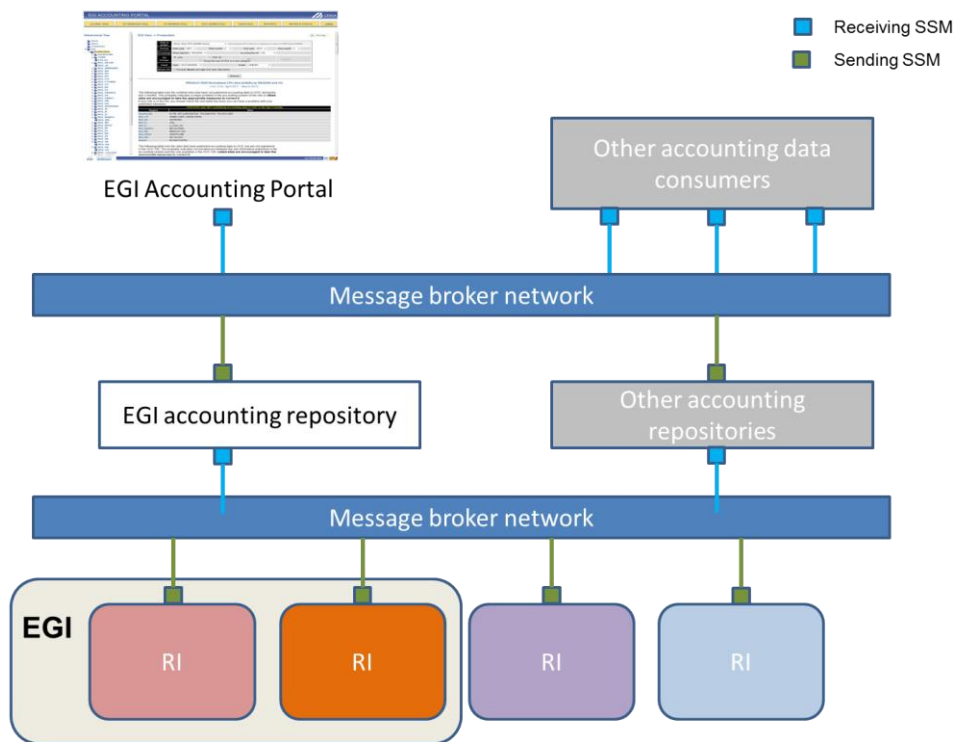
[3] http://myosg.grid.iu.edu/about

[4] http://activemq.apache.org/

Figure 4. Exchange and collection of accounting records through SSM between different RIs.

The EGI accounting infrastructure is a distributed system with central aggregation, where Resource Centres and National Grid Initiatives publish their Usage Records centrally by means of messaging. This infrastructure can be easily extended by plugging more RIs. Both the EGI central accounting repository and the Accounting Portal are consumers of accounting records exchanged between trusted publishers through the message bus.

2.2.2 Federated Cloud support

The Federated Cloud activities within EGI-InSPIRE (task SA2.6) aim at extending the accounting infrastructure of EGI to provide and collect information about usage of cloud resources. Specific parsers are being developed for the cloud middleware being used by the involved RIs. These generate usage records to be stored in the central repository (a test instance of the EGI brokers network used in production is being used for this).

 Once the cloud accounting service will be in production, resource providers will be able to deploy the parser developed for their cloud middleware and deliver the accounting data using the EGI messaging infrastructure. As the cloud usage record will be implemented in the accounting repository, users and resource providers will be able to access the usage statistics in the accounting portal. Aggregated views coupling grid and cloud usage will be implemented as required.

2.3 Service Availability Monitoring

The EGI Monitoring infrastructure is designed to be distributed and scalable, in order to support a dynamic resource infrastructure as EGI. The Service Availability Monitor (SAM) [SAM] system is

based on open source well-established projects such as Nagios[5] and the ActiveMQ messaging system. On top of these open source technologies, custom tools are developed to both meet the EGI requirements and to automate the configuration of every local instance of the distributed monitoring infrastructure.

SAM monitors the services within the production infrastructure and notifies site administrators in case of malfunctioning services and outages that concern their portion of the infrastructure. NGI operations teams and VO operators are also notified for services in their scope. Monitoring results are also used for the calculation of availability and reliability of the infrastructure at different levels of aggregation.

SAM is deployed in a distributed manner. Each NGI (or generally speaking RI) can deploy its own SAM instance that is responsible for monitoring sites and services of that NGI. The central instance deployed at CERN consumes results from all NGI instances, calculates availability and reliability and provides information for the central visualization portal.
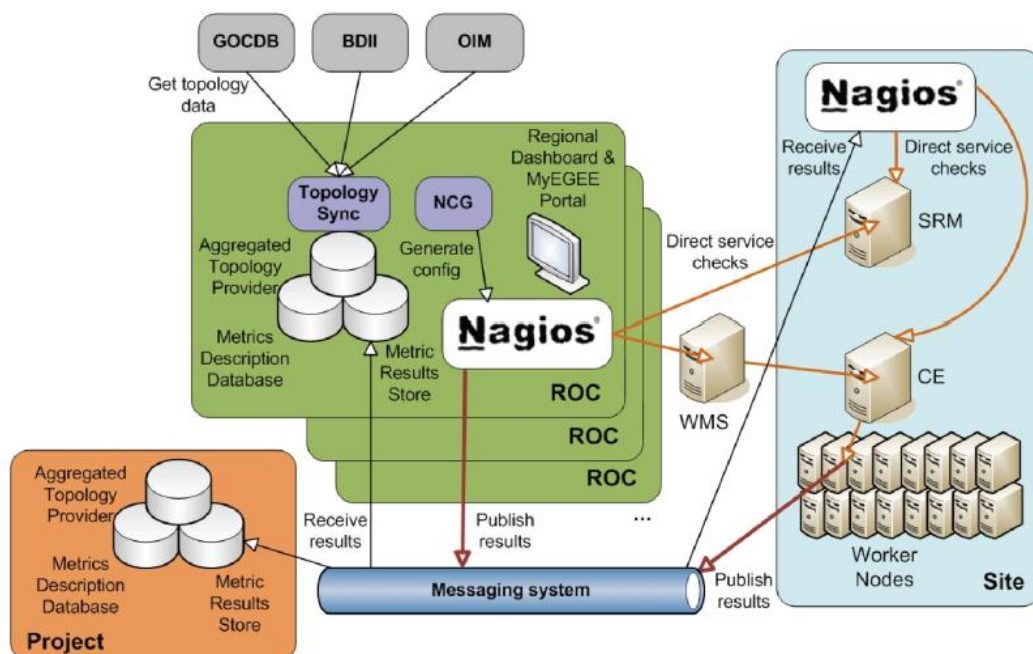

Figure 5. Architecture of the SAM framework

Figure 5 shows the high level architecture of the SAM framework. The list of service endpoints and their topology is fetched by different sources (mainly GOCDB). Nagios is configured merging the topology information with the tests definitions, this allowing for the mapping of the services with the relevant probes. The tests results are stored locally and sent to the EGI metrics result store through the EGI messaging system. The current service status and the historical availability data are available for browsing in the MyEGI portal. In the following paragraph the main components of the system are described.

- **Test execution framework**. The test execution is based on the open source monitoring framework Nagios and the Nagios Configuration Generator (NCG). Nagios is configured by

---

[5] http://www.nagios.org/

merging the topology information with the tests definitions, mapping the services with the relevant probes. The tests results are stored locally and sent to the EGI metrics result store through the EGI messaging system.

- **Databases**. The SAM architecture is built on several databases, deployed both at local and central level. The more relevant data stored in such databases is:

  - Infrastructure topology, the service instances grouped by Resource Centre, gathered from GOCDB and other sources.

  - Profiles, mapping between service types and tests. Profiles are used, for example, to group the tests used to calculate availability and reliability statistics of the services, or the tests that raise alarms in the operations dashboard.

  - Test results and availability and reliability of sites and services.

- **Message bus infrastructure**. EGI deploys a fault tolerant messaging broker infrastructure, based on ActiveMQ, SAM uses the messaging infrastructure for communication between distributed SAM instances, specifically the Nagios boxes deployed at NGI level publish the test results on the message bus, to be consumed by the central instance, which stores the result history data, and by other operational tools.

- **POEM**. Profile Management (POEM) system provides an interfaces and functionality necessary to group different metrics into profiles and, based on those profiles, to configure NAGIOS and all other SAM components. POEM is distributed and various different instances are distinguished using namespaces. One NGI can edit the profiles within their specific namespaces. POEM is hosted together with all the other SAM software in all the SAM instances.

- **Visualization portal.** Although the services status information are available in the Nagios web interface, MyEGI[6] is the main tool used for visualizing the monitoring data. The portal provides different views that summarize the current status of the resource centres in the infrastructure, or the availability historical data of the sites and services monitored. In MyEGI it is possible to select a profile, and show in the different views only the data generated by the tests defined in such profile. The level of detail accessible through the portal spans from site-level availability, to the result of a single test run on a single service. The data available through MyEGI is also accessible through a Programmatic Interface (SAM-PI), which enables other tools (e.g. Operations Portal, VO dashboards) to access test results and availability and reliability of sites and services. The PI allows retrieving other information, i.e. topological data and profiles definitions.

- **Probes.** These are used to test monitored services that are provided by middleware developers and third parties (e.g. NGIs, Nagios community). The probes simulate the user behaviour, in the way they access the services through their public interfaces and test the operations accepted by those interfaces.

2.3.1 Integration of new services in the SAM framework

RIs deploying new types of service, can take advantage of the existing EGI monitoring infrastructure. After the integration of the new services in SAM, RIs can monitor their status together with the current middleware, using the same operational tools, and without adding new workflows.

---

[6] https://grid-monitoring.egi.eu/myegi

To integrate new services with SAM, sensible tests for the service types have to be developed to cover the relevant functionalities defined in the management interface of the product. The probes are subsequently integrated into the SAM Release. Integration also requires from probe developer to provide naming and test configuration (e.g. probe parameters, execution frequency, timeout, etc). The list of currently supported Nagios SAM probes can be found in [PROL]. Additional information about Nagios probe development and integration can be found in [PDI].

2.3.2 Use the SAM infrastructure to monitor non-EGI resources

Custom service types can be registered in GOCDB and tests can be mapped to them creating one ad-hoc profile in the POEM interface, and the service instances can be registered in GOCDB as 'non EGI' scope. Alternatively a dedicated GOCDB service can be used, as explained in section 2.1.1.

Whereas it is possible to raise alarms in the operations dashboard and use them to raise helpdesk tickets, it is not possible to calculate availability statistics including those services, since they are calculated centrally by EGI.

2.3.3 Use the SAM infrastructure to monitor other infrastructures

The SAM framework is tightly built on the topology information provided by GOCDB: to take advantage of the automatic mechanisms for the configuration and visualization of the results RIs need to register their topology in the EGI service catalogue (GOCDB). The current GOCDB release allows to scope sites and services as 'non EGI', and keep a potential new infrastructure operationally decupled from EGI.

It is not currently possible to generate availability statistics for sites not registered as part of the EGI infrastructure.

2.3.4 Federated Cloud support

A set of virtual resources management interfaces and infrastructure services were identified by the Federated Cloud Taskforce as service end-points that need monitoring. In both cases the status of a service end-points impacts the user capability to interact with the cloud infrastructure, i.e. to get access to virtual resources.

The management interfaces are based on open standards: the Open Cloud Computing Interface (OCCI)[7] for virtual machines management, and Cloud Data Management Interface (CDMI)[8] for data management. The infrastructure services are central or distributed operational services needed to manage a distributed infrastructure, currently the Federated Cloud testbed includes the accounting system and the information system.

The extension of the production instance of SAM requires the registration of Federated Cloud service end-points into GOCDB and the integration of Federated Cloud Nagios probes into the SAM distribution. The cloud probes once integrated, will be automatically configured to be run by the NGI Nagios against the cloud service end-points registered in GOCDB. Monitoring results will be centrally stored into the Metrics Result Store (MRS) allowing the computation of monthly availability and reliability statistics.

---

[7] http://occi-wg.org/

[8] http://www.snia.org/cdmi

2.4 Operations Dashboard

In order to operate a distributed infrastructure, management and monitoring information has to be collected and presented in a labour saving way to assist the operators of the infrastructure in their daily work. The dashboard interface[9] combines and harmonizes different static and dynamic information and therewith enables the operators to react on alarms that are notified by the Service Availability Monitoring system, to interact with the sites, to provide first-level support and/or to really operate the Resource Centres by creating and supervising problem tickets on regional as well as central level.

The dashboard allows predefined communication templates and is adaptable to different operational roles (first-level support, regional, central). Resource Centres in the dashboard scope can be regional, central or predefined out of a list and can be sorted and displayed according to numerous criteria to indicate actions needed for a single service, but also for a whole region or even the whole production infrastructure.

The Operations Portal[10] [OPS] content is based on information which is retrieved from several different distributed static and dynamic sources – GOCDB, SAM monitoring system, EGI information system, GGUS, web services, etc. – and gathered into the portal. Interlacing this information has enabled us to display relevant views of static and dynamic information of the EGI production grid. Integrating different technologies and different resources creates high dependencies to the data provided. Consequently, the portal is organized around a web service Lavoisier that provides a transparent integration of each of these resources. The operations portal relies on SAM as source of notifications in case of failed tests (these are received through messaging), and on GOCDB that is the authoritative source for user roles.

2.4.1 Deployment scenarios

The central Operations Portal of EGI provides regionalized views, i.e. pages with information which is selected in order to be relevant to a single site, or Operations Centre of a RI, depending on the user credentials. The scope of the information displayed depends on the user roles in GOCDB.

Alternatively, a local instance can be installed by a RI (the regional package[11] of the operations portal), which synchronizes with the central one. The regionalized instance only includes a subset of modules, i.e. those that are relevant at a RI level. Other Operations Portal modules (such as the VO dashboard, the VO registration module, the broadcast tool etc.) are only deployed centrally.

The application is composed of:

- a web service named Lavoisier configured to handle Nagios notifications, store and provide data cache from GOCDB, GGUS and to generate metrics reports about the use of the dashboard (alarms raised, tickets handled etc.).
- a PHP web application to provide a user interface.

The regional instance is linked with the central instance of Lavoisier: creation, update, delete of records are synchronized so as not to disrupt central oversight operations. Synchronisation is achieved

---

[9] https://operations-portal.egi.eu/dashboard

[10] http://operations-portal.egi.eu/

[11] https://svn.in2p3.fr/operations-portal/regional/

---

through REST and SOAP and records are synchronized every 5 min by using php scripts. Any problem detected during the synchronisation is reported in a mail sent to webmasters.

The central and the regional instances have been built on the same model to behave in the same way and to be easily interoperable. As shown in Figure 6, the architectures of the central Operations Portal and the regional one are exactly the same. The two types of instances only differ in their configuration.
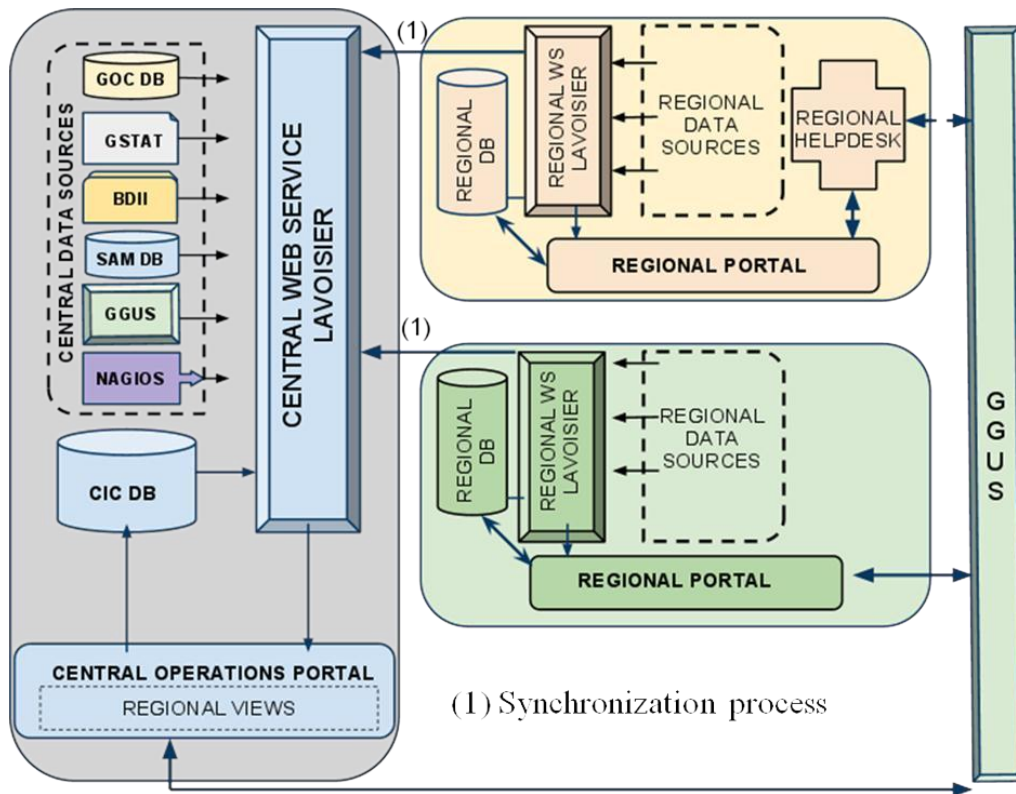

Figure 6. Architecture of the regional and central Operations Portal.

2.4.2 Federated Cloud support

The Operations Dashboard provided by the Operations Portal is software agnostic and only depends on the availability of notifications generated by SAM and distributed through the message bus.

A full integration with the EGI monitoring system will allow resource providers to use the Operations Dashboard to have an overview of the Nagios alarms for cloud resources as well as grid resources.

2.5 Information discovery

Information discovery is a capability that helps find the required resources that have been registered with it within the production infrastructure [D5.2]. The information collected about such resources is made available through well-known instances that provide the data to some logical collection, infrastructure wide, regional, site, domain, etc.

The information discovery service provides information to support the following use cases.

- Service Discovery (SD): information which describes what services exist and what their properties are. This is expected to be information which changes rather slowly and hence may usefully be cached. Such information will typically be used to filter the results of a query, e.g. in the WMS JDL Requirements expression, or be returned as service metadata.

- Service Selection (SS): information, often highly dynamic, which facilitates the choice of a particular service at a given time among several which satisfy the SD requirements. This may for example reflect the current load on the service. This will typically be used to rank the results of an SD query, e.g. in the Job Description Language Rank expressions used by workload schedulers.

- Monitoring (M): dynamic or static information which may be useful to monitor and display the properties and state of services across the Grid. This may be presented either individually for each service or aggregated across a set of sites. Specific monitoring tools are likely to display only a subset of this information.

- Oversight (O): information, generally static or slowly changing, which can be aggregated to provide a high-level picture of Grid resources, for example the total installed CPU and disk capacity. This category also includes information which is useful for the overall management of the Grid, for example the distribution of versions of particular services.

- Diagnostic (D): information which facilitates the tracing of problems with Grid services or within the information system itself

Clients to such service must be able to search, filter, and order the available information until their initial request is satisfied. The LDAPv3 (RFC 4530) protocol and search syntax is used to query information from the information discovery services and to encapsulate the information payload relating to the services being offered within the production infrastructure that is exchanged between instances.

EGI adopts the GLUE specification as an information model for Grid entities described using the natural language and UML Class Diagrams. GLUE, as a conceptual model, is designed to be independent from the concrete data models adopted for its implementation. EGI is currently based on version 1.3 of the specification [GLUE1], and is progressively adopting version 2.0 [GLUE 2.0]. An EGI GLUE 2.0 profile was drafted, proposing the usage of the GLUE 2.0 schema in EGI. It extends the schema specification document [GLUE2] with more detailed semantics for particular attributes, specifies conditions under which objects and attributes should and should not be published, and defines some additional information to be published. It also includes guidelines for validating the accuracy of the published information.

All deployed software stacks (e.g. EMI and IGE) are publishing or a planning to publish information into a GLUE-based information discovery service. This makes resources offered discoverable regardless of the specific interface that is exposed to gain access to them.

2.5.1 Federated Cloud support

Resources of the EGI federated cloud testbed publish information to a LDAP servers using an extended GLUE 2.0 schema. The information published includes: the endpoints of the management interfaces (web services and web sites), the amount of resources provided by the resource provider and the virtual images available in the site to be instantiated by the cloud users.

The proof of concept of a cloud information system uses the same open standards adopted by the grid equivalent: GLUE schema and LDAP. No use cases have been collected that require a common information system providing information for grid and cloud, although this would be technically possible and probably desirable to reduce deployment costs.

An extension to the GLUE 2 schema will be likely proposed; if this proposal will be endorsed by the GLUE working group at OGF, an integrated information system for grid and cloud will become possible.

While the BDII services publish a dynamic summary of what is online, the EGI services are registered in GOCDB to keep a static picture of which services the resource providers are deploying, the scheduled downtimes, the resource centres topology and the relevant contacts of the site staff people (see section 2.1.4).

3 COLLABORATION PLATFORM

This chapter provides information on how external RIs can deploy the EGI helpdesk, which is part of the EGI Collaboration Platform and of the operations service portfolio. Various integration scenarios are illustrated in the following sections.

3.1 EGI Helpdesk

The user support infrastructure in use within the EGI Helpdesk is distributed consisting of various topical and regional helpdesk systems that are linked together through a central integration platform, the GGUS helpdesk [GGUS]. This central helpdesk enables formalized communication between the submitter of the incident record and all partners involved in user support by providing an interface to which all other tools can connect and enabling central tracking of a problem, independent of the origin of the problem and the tool in which the work on the problem is done.

The interlinking of all ticket systems in place within EGI enables to pass incident records from one system to the other in a way that is transparent to the user. By exposing agreed interfaces, a hierarchy of tree of interworking helpdesk systems operated by different Ris, can be implemented allowing for transparently exchanging incident records across different Ris. It also enables the communication and ticket assignment between experts from different areas (e.g. middleware experts and application experts) while at the same time allowing them to work with the tools they are used to, as well as the assignment of incidents to the affected RI.

Beside of the central helpdesk system GGUS, EGI also provides the xGUS helpdesk[12] template which has been developed for Ris and user communities who want to build up an own user support infrastructure that is seamlessly integrated with the one operated by the Ris. It contains all basic helpdesk functionalities including user administration and certificate access. Portal administrators from the client RI or user community can customize the portal to their specific needs. Via web, they can edit, add and delete the support units, variables which are used for the classification of tickets like 'Type of problem', 'VO' etc. and the hyperlinks to related web pages displayed on the portal. The use of xGUS guarantees a consistent user support infrastructure throughout the whole EGI project.

In the following sections we present various RI helpdesk integration scenarios that suit different scenarios according to the user communities supported and the technologies deployed by the Ris.

3.1.1 User-specific helpdesk front-ends

A number of European Ris complement each other and meet different needs of the same user communities. The combination of EGI, PRACE and EUDAT can serve as an example of this. A scientific community might use the different computing infrastructures in parallel for different computing use cases and store their data in the distributed data infrastructure.

To support the coupled usage of various Ris by scientifc users, a strong collaboration between Ris is needed. If for exapmple each of the Ris offered different helpdesks to the users to report problems with this specific infrastructure, this might discurage the combined usage of the various infrastructures, since the user in an ideal case wouldn't even have to know which part he was currently using, as this would be hidden by the virtual research environment he is using.

---

[12] http://xgus.scc.kit.edu/

To make the user support work properly in such an environment of federated e-Infrastructures, the support tools – the heldesk systems – need to be federated as well.

The main requirement for the seamless federated support across multiple Ris is that the helpdesk systems are interfaced with each other. Doing this as an N-to-N solution is not scaleable, therefore there has to be some kind of integration platform to which all helpdesks involved are interfaced and which handles the transferes of tickets between the systems. GGUS – the helpdesk system of EGI – could serve as this platform.

There are two different approaches to the presentation of this system of federated helpdesks to the user. One option (Figure 7) would be to have one single front-end used by all scientific communities using the federated e-Infrastructures. This front-end then would have to cater for all the use cases problems might be reported about, which could lead to a rather overloaded and difficult to use portal.
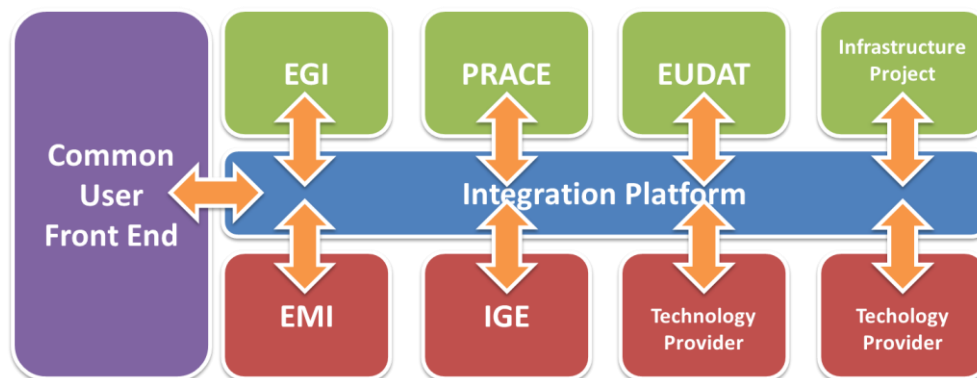

Figure 7. Example of integration between helpdesk systems of different Ris based on a shared front-end system (EGI, EUDAT, PRACE, etc.), of technology providers offering 3rd level support (e.g. the European Middleware Initiative – EMI, the Initiative for Globus in Europe – IGE, etc.)

The second option (Figure 8) would be to include community specific helpdesk front-ends in the virtual research environments used by the communities. Both of these approaches would hide the complexity of the federated helpdesks from the users. The second appraoch would also enable a community specific first line support to be performed to filter out community specific problems before they reach the infrastructure helpdesks. The xGUS system can serve as a community specific front-end for interested communities.
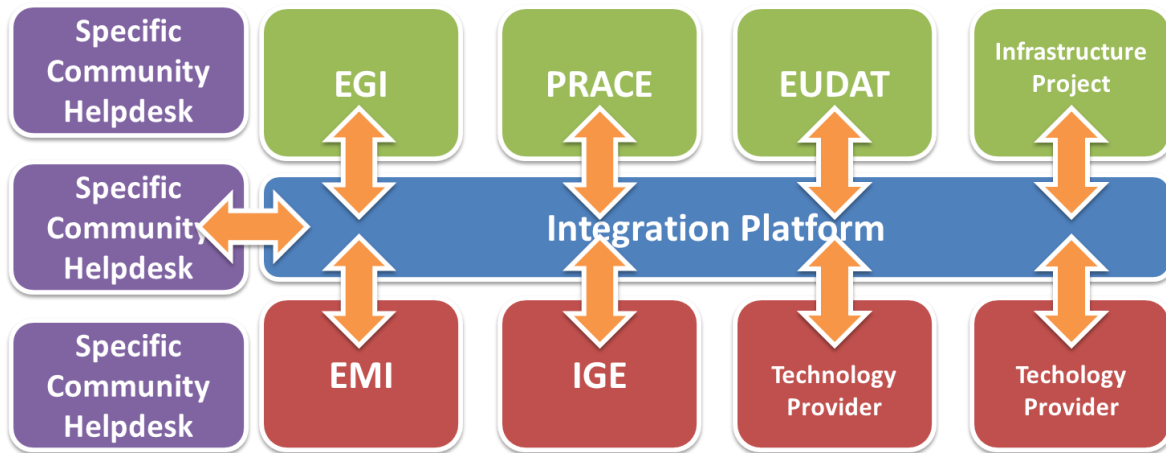
Figure 8. Example of helpdesk integration scenario based on the provisioning of user-community specific front-ends interfaced to a tree of interoperable RI helpdesk systems.

3.1.2 Shared helpdesk support units

Ris can opt for a dedicated front-end for their users (xGUS) that is interfaced with a common set of Sus hosted by the GGUS backend (Figure 9). This solution is suitable in case of user communities whose resource usage is confined within a given RI (each user community is exposed to a dedicated helpdesk front-end), where however Ris deploy a shared set of software stacks (e.g. UNICORE in EGI and PRACE). This solution allows different Ris to share a single interface towards the Technology Providers of interest, and at the same time allows for consolidating support across different user communities of the same technology.

In this scenario each user communities can decide to use standalone dedicated xGUS instance operated by EGI. In this option, information about incidents handled in one users helpdesk is not propagated to other systems and for this reason may cause the situation where one common incident is handled in multiple systems separately.
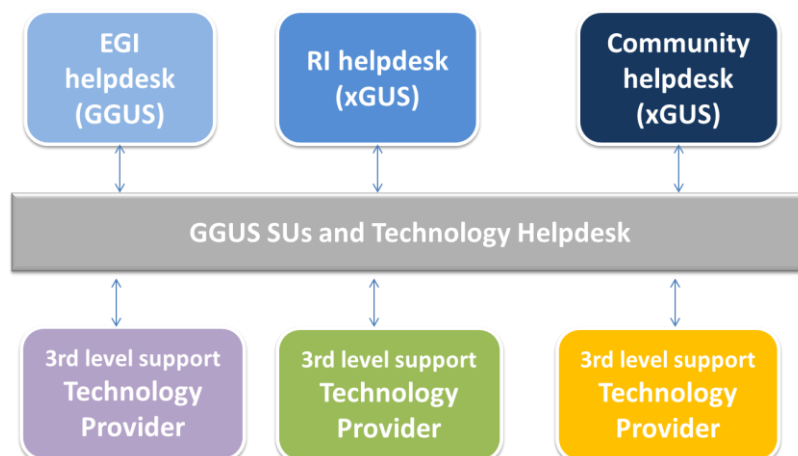

Figure 9. Different RI helpdesk front-ends sharing the same GGUS backend and the 3rd level support provided by Technology Providers.

3.1.3 Different RI helpdesks

A standalone xGUS system can be provided to a user community and can be interfaced with external RI helpdesk systems (Figure 10). In this scenario incident records originated in the xGUS instance – operated by EGI – can be propagated to other RI support systems, if the incident is specific to a given RI. Information is synchronized across the different helpdesk systems so that changes to tickets that are modified locally in a RI helpdesk, are propagated upstream to the xGUS instance.

This solution prevents de-synchronization of support teams while allowing Ris to adopt an own helpdesk system, and is the integration scenario that is being currently investigated by EGI and PRACE. The deployment of a community-specific helpdesk requires a user community to internally provide a team of experts who are responsible of ticket triage, i.e. of solving the problem or of assigning the records to a specific RI as deemed appropriate.
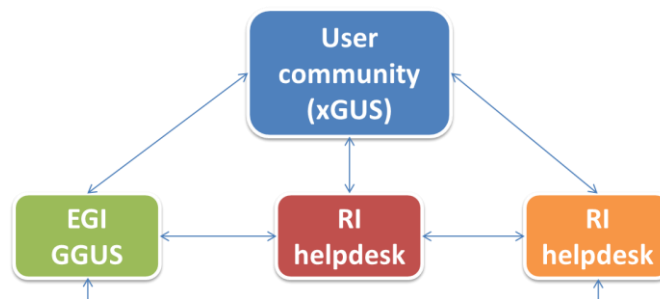

Figure 10. Different RI helpdesks are being operated with the capability of exchanging incident records between them to make sure the problem is handled by the appropriate RI. A fully integrated xGUS dedicated helpdesk system can be provided to a user community with the capability of exchanging tickets with each RI helpdesk system.

4 CONCLUSIONS

The adoption of EGI operations tools, deployed and developed within EGI-InSPIRE, allows the reuse of existing solutions that address the typical needs of any distributed RI. These needs are typically resource and technology independent.

EGI can share these building blocks in a variety of configuration options. EGI can operate these tools in a software-as-a-service fashion for other RIs, can service RIs through the existing production instances of EGI, or alternatively provide the tools as open-source software for the local deployment of independent instances that can be easily integrated with existing EGI production services.

5   REFERENCES

APEL	Roger, W., Jiang, M., Del Cano, C. et alt.; An APEL Tool Based CPU Usage Accounting Infrastructure for Large Scale Computing Grids, in the Proc. Of Int. Symposium on Grid Computing 2010, Taipei (TW), Mar 2010 (http://event.twgrid.org/isgc2010/abstract/AnAPELToolBasedCPUUsageAccountingInfrastructureforLargeScaleComputingGrids.pdf)
CAR	Definition of the Compute Accounting Record, EMI Technical Note (https://twiki.cern.ch/twiki/bin/view/EMI/ComputeAccounting)
D4.3	Ferrari, T.; EGI Operations Architecture: Grid Service Management Best Practices, EGI-InSPIRE Deliverable D4.3, Dec 2011 (https://documents.egi.eu/document/763)
D5.2	UMD Roadmap, EGI-InSPIRE Deliverable D5.2, Mar 2011 (https://documents.egi.eu/document/272)
GGUS	Global Grid User Support (GGUS) (https://wiki.egi.eu/wiki/GGUS)
GOCDB	Grid Configuration Database (GOCDB) (https://wiki.egi.eu/wiki/GOCDB)
EAP	EGI Accounting Profile (https://indico.egi.eu/indico/conferenceDisplay.py?confId=723)
EGP	EGI Profile for the Use of GLUE 2.0 Information Schema, EGI Draft, Aug 2012 (https://documents.egi.eu/document/1324)
GLUE1	Andreozzi, S.; Burke, S.; Field. L. et alt.; GLUE Schema Specification version 1.3 (https://forge.ogf.org/sf/docman/do/downloadDocument/projects.glue-wg/docman.root.background.specifications/doc14185), Jan 2007
GLUE2	Andreozzi, S.; Burke, S.; Ehm, F. et alt.; GLUE Specification v. 2.0, OGF Recommendation GFD-R-P.147, Mar 2009 (http://ogf.org/documents/GFD.147.pdf)
M506	EGI Platform Roadmap, EGI-InSPIRE Milestone MS506, April 2012 (https://documents.egi.eu/document/970)
OPS	EGI Operations Portal (https://wiki.egi.eu/wiki/Operations_Portal)
PDI	Nagios Probe Development Guidelines (https://tomtools.cern.ch/confluence/display/SAMDOC/Probes+Development)
PROL	SAM Probes: https://wiki.egi.eu/wiki/SAM#Tests_and_probes
SAM	Service Availability Monitoring (https://wiki.egi.eu/wiki/SAM)
SSM	Secure Stomp Messenger (https://wiki.egi.eu/wiki/APEL/SSM)
STAR	Nilsen, J. K. et alt.; EMI StAR – Definition of a Storage Accounting Record, OGF draft, Sep 2012