EGI-InSPIRE

Integrating Resources into the EGI Production Infrastructure

EU MILESTONE: MS414

Document identifier:	EGI-MS414-V0-17
Date:	**28/08/2011**
Activity:	**SA1**
Lead Partner:	**KTH**
Document Status:	**DRAFT**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/650

Abstract

This document describes and defines the operational interfaces that must be supported for resources to be integrated into EGI. This includes operational tools provided by the EGI-InSPIRE JRA1 activity and procedures and policies defined to ensure interoperability within EGI and in the interaction with other DCIs, the adoption of best practices and compliance with service level agreements.

II. DELIVERY SLIP

		Name	**Partner/Activity**	**Date**
From		Michaela Barth	KTH/SA1	
Reviewed by		**Moderator:** **Reviewers:** <<To be completed by project office on submission to AMB/PMB>>		
Approved by		**AMB & PMB** <<To be completed by project office on submission to EC>>		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
0	05/07/2011	Incomplete placeholder	Michaela Barth/KTH
1	06/07/2011	ToC	Michaela Barth/KTH
2	12/07/2011	Input on Operations Portal	Cyril L'orphelin/IN2P3
3	19/07/2011 27/07/2011 29/07/2011 01/08/2011	Input on GOCDB, Input on ARGUS in general, ARGUS and gLite Input on GOCDB Input on ARGUS and ARC Comments	David Meridith/STFC Alvaro Simon/CESGA Torsten Antoni/KIT Ali Gholami/Nordugrid Michaela Barth/KTH
4	02/08/2011	Input on ARGUS and UNICORE	Krzysztof Benedyczak/ UWAR
5	02/08/2011	Some corrections	
6	04/08/2011	Input on Accounting	John Gordon/ STFC
7	04/08/2011	Going through general parts	Michaela Barth/KTH
8	05/08/2011	Going through accounting input Input on Monitoring	Cristina del Cano Novales/ STFC Emir Imamagic/SRCE
10	12/08/2011	Review	T. Ferrari/EGI.eu
11	21/08/2011	Added references	Gert Svensson/KTH
12	21/08/2011	Added some text on Argus	Alvaro Simon/CESGA
13	22/08/2011	Addressing review comments	Michaela Barth /KTH
14	23/08/2011	Minor cleanup and addressing some	Gert Svensson/KTH

15	24/08/2011	comments from Oscar Koeroo	
15	24/08/2011	Adding text from Oscar Koeroo and adding changes from Ali Gholami	Gert Svensson/KTH
16	26/08/2011	Addressing review comments	Emir Imamagic /SRCE
17	28/08/2011	Addressing review comments New figure from Alvaro Simon	Gert Svensson/KTH

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

Various term definitions are available in the EGI Glossary at: https://wiki.egi.eu/wiki/Glossary; acronyms are defined at http://www.egi.eu/about/glossary/.

VII. PROJECT SUMMARY

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

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

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

The objectives of the project are:

1. The continued operation and expansion of today's production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.
2. The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
3. The support for current heavy users of the infrastructure in earth science, astronomy and astrophysics, fusion, computational chemistry and materials science technology, life sciences and high energy physics as they move to sustainable support models for their own communities.
4. Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
5. Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure, so as to provide transparent access to all authorised users.
6. Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids) and heterogeneous resources (e.g. HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

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

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

VIII. EXECUTIVE SUMMARY

This document defines and describes the operational interfaces that must be supported for resources to be integrated into the European Grid Infrastructure (EGI), and is an updated version of MS407 [R 6].

During the first year of the project activities focussed on the integration of four middleware stacks: ARC, gLite, GLOBUS and UNICORE. The integration of ARC was completed during the first project year, while two task forces were set-up to steer the integration of UNICORE [R 73] and GLOBUS[R 74] which saw the involvement of the technology providers and of the relevant Resource Providers. This document presents the works of the task forces and accomplishments.

For each of the operational tools we describe the steps necessary to integrate a new middleware stack into the production infrastructure. This is followed by a detailed analysis of each middleware stack and the related medium-term development plans relevant to their operational interoperability.

In summary the integration of gLite and ARC can be considered complete while for UNICORE and Globus the now existing SAM Nagios probes have to be fully integrated and an operational subset of them has to be defined. Also integration into the accounting system is still outstanding for those two middlewares. The integration into the Operational Dashboard is expected to work out of the box, though (see also integration overview table [Table 1]).

This document also gives an overview of the status of EGI operational procedures and policies needed for the integration of new resources.

Finally, we discuss further integration requirements coming from different sources, like NGIs, other DCIs and above all from our successful integration task forces, and conclude with our future plans around the completion of the integration of UNICORE and Globus, and on the integration with desktop Grids and PRACE.

TABLE OF CONTENTS

1 INTRODUCTION

In order to add new resources into the EGI production infrastructure a basic set of operational interfaces that must be supported by the new resources, has to be defined and described in their basic functionality.

Different resources will use different middleware components. EGI-InSPIRE will support the Unified Middleware Distribution (UMD) for deployment on the production infrastructure, which integrates software for multiple technology providers.

Operational tools such as the GOC Database (GOCDB) and the SAM/Nagios monitoring tools are key software components for a reliable and stable operation and monitoring of the infrastructure. The current operational tools may change in the future, but they provide the starting point when comparing the operational interoperability of different middleware components for each of the operational tools currently in use.

Operational procedures and policies are needed to enforce the application of the agreed basic set of operational interfaces to be supported by all resources. The EGI procedures and policies have been adapted and new requirements were identified which turned into new procedures and policies relevant for the integration of new resources.

2   TECHNOLOGY AND OPERATIONAL TOOLS

The EGI-InSPIRE project continues to evolve the blueprint on how to successfully run a federated European Grid Infrastructure. A certain amount of rationalization and optimization is necessary to pick up best practice within the community and to create a sustainable model for operating a growing pan-European grid infrastructure that builds on nationally and regionally funded grid initiatives who want to work together.

Availability and reliability measurement, registration of services, information indexing, monitoring, accounting, user and operational support in EGI currently rely on operational tools which are developed in EGI-InSPIRE JRA1 [R 2].

While different middleware stacks are supported by EGI for deployment in the resource centres, the central and distributed instances of the operational tools are operated by a small number of partners committed to provide such services for National or Regional Grid Initiatives, or even for the whole EGI.

EGI will need to deploy several middleware stacks according to the requirements of users and site managers. Presently, gLite and ARC can be viewed as fully integrated into all the operational tools, whilst some smaller adaptations are still needed due to changed and more standardized interfaces of the operational tools enabling broader access to other types of middleware. Globus and UNICORE operational integration is in full progress also thanks to the specialised integration task forces. The comprehensive integration is a short-term objective of the first phase of the project.

In a second phase new types of resource will be integrated, such as virtualization, digital libraries and repositories, desktop grids, High Performance Computing, etc.

2.1   Interfaces

The basic operational interfaces that must be supported for resources to be integrated into EGI consist of a management interface, a monitoring interface, an accounting interface, a support interface and an additional graphical dashboard interface which collects and presents the information provided by the others and ties them together in a meaningful way to facilitate daily oversight grid monitoring duties.

MANAGEMENT INTERFACE.

An important operational interface of a resource is the capability to be put in downtime if under maintenance, the capability to undergo a certification process and thereby reach production status, and the capability to be monitored to assess its operational security level. GOCDB is the tool of choice for fulfilling these management tasks. It provides information about the services, where and who to contact on a management and technical level, as well as in case of security issues. A first step towards integration of resources is therefore the possibility to enable the registration of new types of services provided by these resources in GOCDB.

MONITORING INTERFACE.

The next step is to describe and advertise the resources using the OGF GLUE standard schema (GLUE 1.3 and 2.0) [R 14]. This enables the construction of a unified topology which is necessary for the monitoring of the infrastructure. One possible monitoring tool fulfilling these requirements is for example Nagios, which allows all relevant services to be probed at regular intervals to assess their operation. Such a test execution and notification environment is needed for the fast identification and consequently fast resolution of functional problems that affect the infrastructure. General monitoring of services is also needed to produce the results that are then consumed to produce availability and reliability reports.

ACCOUNTING INTERFACE
Accounting is about the collection of resource usage information. Usage information can be collected for various resource types, however the current accounting technology only allows the accounting of compute resources. The accounting infrastructure currently comprises a central repository that collects information from the individual Resource Centres and/or grid infrastructures. Usage records are exchanged among the publishers and consumers by means of a message passing infrastructure.

SUPPORT INTERFACE
The grid technology that is deployed and integrated needs to be supported in case of installation, configuration and functionality issues. EGI provides first and second-level support, while specialized support is typically offered by the technology providers themselves.
The EGI Helpdesk (GGUS) is a distributed support system with central coordination. The EGI Helpdesk is a common infrastructure to exchange trouble tickets between different support units.

OPERATIONS DASHBOARD INTERFACE
Failures that are detected by the monitoring system generate notifications that produce alarms in the Operations Portal. In case of alarms, Resource Centre administrators are contacted by submitting trouble tickets via the EGI Helpdesk.

USER MANAGEMENT INTERFACE
Although not explicitly being an operational tool per se, user membership management interfaces are necessary for authentication and authorization. These capabilities influence the work with all the other operational tools.

In the following sections the various integration interfaces are illustrated in detail.

2.2 Overview Status of Middleware Integration for each Operational Tool

During the first year of the project activities were focused on the integration of various technologies: ARC, gLite, GLOBUS and UNICORE. The current integration status is summarized in Table 1.

Table 1. Status of integration of ARC, gLite, GLOBUS and UNICORE

	gLite	ARC	UNICORE	GLOBUS
GOCDB	Completed	Completed	Completed (first services are being registered)	Completed (first services are being registered)
Monitoring 1. **Nagios probes written,** 2. **Probes integrated,** 3. **Definition of an OPERATIONAL set for integration into the operations dashboard**	Completed	Completed (integration into SAM release 7)	Probes have been written, integration foreseen for SAM release 13	Probes have been written and will be supported by IGE in the future, integrated in SAM release 11, definition of an OPERATIONAL probes requires certified sites
Operational Dashboard	Completed	Completed	To be done (should work automatically after definition of an operational set of Nagios probes)	To be done (should work automatically after definition of an operational set of Nagios probes)
Accounting	Completed	Completed	In progress	In progress
3rd level support in GGUS (access to expert teams via the Deployed Middleware Support Unit)	Completed	Completed	Completed	Completed

2.3 *Management Interface*

2.3.1 Functionality

A management interface allows Resource Centres to store, maintain and view the topology of the production infrastructure and the basic information about the respective resources within it. Such an EGI management interface contains information about:

- 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 access points for these resources.
- Participating people and their roles within EGI operations.

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. The presented information can be a combined view of different regionalized or otherwise separated instances with their own local inputs.

2.3.2 Requirements

The EGI management interface has to support the functionality described above. System and security contacts and higher level organizational management contacts for a Resource Centre need to be easily identified. The management interface may provide finer granularity for contact details by marking extended expertise on a specific middleware stack or an affinity to certain types of service(s).

Additionally, it must be possible to register new kinds of service types, groups or sites within the management interface. A site should be able to contain services from different middleware stacks. The description and/or the name of the service type should also contain information about the respective technology provider.

Such a database needs a role-based interaction model, so that people responsible for certain Resource Centres, services or resources can update and maintain the various entries representing the entities under their responsibility within typical daily operations scenarios. In particular, basic service status information shall be easily viewable and changeable. It shall be easily possible to register a service of a known service type, to edit system administration information and put whole sites or single resources in and out of downtime according to predefined procedures. It shall be easy to identify whether a resource is monitored or not by the corresponding monitoring system. This monitoring bit can be set separately or implicitly within the different production states.

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

Since the management interface provides much needed basic information on the topology of the production infrastructure and its contact points, we expect a plug-in to an approved dashboard interface to be in existence or easily implementable by using canonical standards. Even though the information is mostly static, a regionalized version with a central collecting portal of the management interface would of course be preferred in order to emphasize the distributed nature of the grid community, to avoid single points of failure and to manage local resources that are not part of EGI.

2.3.3 Integration into GOCDB

Services registered in GOCDB have; 1) a 'Service Type' identifier, 2) a required 'Service Endpoint' instance and 3) an optional 'Endpoint Location'.

1. **Service Type:** a unique name that identifies a type of software component deployed on a Grid, including middleware (e.g. CE, WMS, SRM) and/or operational components (e.g. MessageBroker, RegionalNagios). The naming scheme for new service types follow a reverse DNS style syntax, usually naming the technology provider followed by technology type, i.e. '<provider>.<type>' (e.g. 'unicore6.StorageFactory'). This is consistent with the proposed EMI service registry naming scheme from GLUE2.0 that defines a service type enumeration. It would be preferable to rename all existing service types using this scheme, but this is potentially problematic for existing services that depend on established legacy names. The current list of service type definitions are given in [R 51].

2. **Service Endpoint:** represents a deployed instance of a service type.

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

2.3.3.1 Procedure for registering new Service Types

New service types can be registered by GOCDB administrators. Once registered in GOCDB, users (site administrators, regional managers) can declare instances of the new service type as required. The complete procedure to integrate new service types is as follows.

1. If the service type is already registered in GOCDB, service endpoints can be added by users of GOCDB following the established procedure.
2. If the service type is not registered, a request should be made to the OTAG through the respective Resource Provider in the RT system for its inclusion in GOCDB. If the new service type belongs to a previously undeclared middleware stack, then a strategic decision is required to ensure only officially supported middleware is integrated into GOCDB. If the request is approved, it is communicated to the GOCDB developers to add the new service type.
3. The requesting party is notified (either the request is rejected or completed).

2.3.3.2 Regular review of the list of available service types

A regular review of the supported GOCDB service types will be made. This is the responsibility of GOCDB developers, who will consult the Technology Collaboration Board (TCB) together with the Operations Management Board (OMB).

2.3.3.3 Integrated operational tool types

- **egi.APELRepository**: [Central service] The central APEL repository
- **egi.AccountingPortal**: [Central service] The central accounting portal
- **egi.GGUS**: [Central service] The central GGUS
- **egi.GOCDB**: [Central service] The central GOCDB
- **egi.MSGBroker**: [Central service] The central message broker
- **egi.MetricsPortal**: [Central service] The central metrics portal
- **egi.NetworkPortal**: [Central service] The central network portal
- **egi.OpsPortal**: [Central service] The central operations portal
- **egi.SAM**: [Central service] The central SAM monitoring
- **egi.GRIDVIEW**: [Central service] The central gridview portal
- **egi.GSTAT**: [Central service] The central GStat portal
- **ngi.SAM**: [Regional Service] NGI-level Nagios monitoring box
- **ngi.OpsPortal**: [Regional service] NGI-level regional operations portal instance

The following service types are obsolete and will be removed:

- **OpsTool**: [Central service] generic service representing an operation tool (topology repository, dashboard, helpdesk system...) – to be replaced with individual service types listed above
- **MSG-Broker**: [Central service] A broker for the central/backbone messaging system – to be replaced with egi.MSGBroker
- **Site-NAGIOS**: [Site service] site-level Nagios monitoring box – to be removed
- **Regional-NAGIOS**: [Regional Service] ROC-level Nagios monitoring box – to be replaced with ngi.SAM
- **Project-NAGIOS**: [Central Service] project-level Nagios monitoring box – to be replaced with egi.SAM

2.3.3.4 Integrated gLite service types

- **CE**: [Site service] The LCG Compute Element. Currently the standards CE within the gLite middleware stack. Soon to be replaced by the CREAM CE.
- **CREAM-CE**: [Site service] The CREAM Compute Element is the new CE within the gLite middleware stack.
- **APEL**: [Site service] This is a "dummy" Service Type to enable the monitoring tests for APEL accounting. All sites must have one instance of this Service Type, associated with a CE.
- **UI**: [User service] The User Interface. Can be installed by users but more commonly installed by a site.
- **SRM**: [Site service] Storage Resource Manager. Mandatory for all sites running an SRM enabled storage element.
- **Central-LFC**: [Central service] An instance of the gLite file catalogue which holds entries for all files owned by a particular VO. NOTE: An LFC can be both Central and Local.
- **Local-LFC**: [Site service] An instance of the gLite file catalogue which holds entries for files owned by a particular VO, at your site. NOTE: An LFC can be both Central and Local.
- **WMS**: [Central service] gLite Workload Management Service. Acts as the broker for matching user jobs to available computing resources.
- **VOMS**: [Central service] VO Management System [R 9]. Part of the authentication and authorization system. This service only needs to be installed on the request of a VO.
- **MyProxy**: [Central service] The My Proxy service is part of the authentication and authorization system. Often installed by sites installing the WMS service [R 15].
- **LB**: [Central service] gLite Logging and Bookkeeping. Usually installed by sites running a WMS. One LB service can support several WMS instances.
- **Site-BDII**: [Site service] This service collects and publishes site's data for the Information System. All sites MUST install one Site-BDII.
- **Top-BDII**: [Central service] This is the "top-level BDII". These collect data from site-BDIIs and publish the data. Only a few instances per region are required.
- **AMGA**: [Central service] gLite metadata catalogue. This service only needs to be installed on the request of a VO.
- **FTM**: [Site service] gLite File Transfer Monitor. Monitors the FTS service at a site.
- **FTS**: [Central service] The gLite File Transfer Service manages the transfer of files between sites. This service only needs to be installed on the request of a VO.
- **VO-box**: [Site service] The gLite VO box allows a VO to run their own services at a site. This service only needs to be installed on the request of a VO.
- **gLite-APEL**: [Site service] The gLite-APEL hosts the site Accounting client (3.2 replacement of the MonBox)
- **gLExec**: [Site service] A light-weight gatekeeper to authenticate and authorize credentials according to local site policy and execute commands.

2.3.3.5 Integrated ARC service types

As of release 0.8 of ARC, the ARC-CE runs a resource BDII with GLUE schema 1.3, in the same way as gLite resources. Hence setting up a special site BDII is no longer needed. More details are found in [R 22].

- **ARC-CE**: [Site service] The Compute Element within the ARC middleware stack.
- **SGAS**: [Site service] An accounting service used by ARC.

2.3.3.6 Integrated UNICORE service types

- **unicore6.Registry**: [Central service] All UNICORE services register here; clients ask the registry for available services in the Grid. Normally one Registry per Grid infrastructure which collects URLs of services.
- **unicore6.Gateway**: [Site service] Sits in front of one or more UNICORE services as a gateway to the internet. Normally one Gateway per site.
- **unicore6.TargetSystemFactory** [Site service] used as an entry-point for submitting single jobs. It can create Target System Services (TSSs) and submit jobs to those TSSs.
- **unicore6.StorageFactory** [Site service] Creates StorageManagement instances. A user can create dynamic storage management services for own purposes with it. Often used to provide file space during workflow execution.
- **unicore6.StorageManagement** [Site service] Provides an abstract file system-like view on a storage resource. A Storage Management Service (SMS) can be created by a Storage Factory or can be configured statically way by a configuration file.
- **unicore6.ServiceOrchestrator** [Site service] Handles dispatching of a workflow's atomic jobs, and brokering. Normally there is one per grid infrastructure.
- **unicore6.WorkflowFactory** [Site service] Used as an entry point for submitting workflow jobs. The Workflow factory is creating workflow instances and can submit workflows to them. It is the workflow submission equivalent to the Target System Factory used for single job submission.
- **unicore6.UVOSAssertionQueryService** [Site service] Provides data and user information via the SAML standard as needed for authorization and environment customization.

unicore6.UNICOREX was removed from the list as decided by the UNICORE integration taskforce, since it is a just an undefined collection of other services like TargetSystemFactory, StorageFactory, StorageManagement.

Other services which might need to be integrated in the future include:
unicore6.BESFactoryPortType,
unicore6.TSI (unlikely, since only internal),
unicore6.locationManager (unlikely since only for backwards compatibility) ,
unicore6.Catalogue,
unicore6.GridResourceInformationService (unlikely mostly coupled to the service orchestrator),
unicore6.TracerService (unlikely since strongly coupled to the workflow service),
unicore6.CISRegistryPortType,
unicore6.CISInfoProvider,
unicore6.XUUDBQuery (internal, but might give some benefits),
unicore6.GridBeanService.

2.3.3.7 Integration of Globus resources

- **GRAM5**: [Site service] job submission service for Globus version 5.x (GRAM5)
- **globus-GRIDFTP**: [Site service] storage endpoint and data transfer service for the Globus middleware stack
- **globus-GSISSHD**: [Site service] certificate based interactive login service for the Globus middleware stack

2.4 Monitoring Interface

2.4.1 Functionality

A monitoring interface monitors the resources presented within EGI to ensure the infrastructure's reliability and to quickly find causes of failure.

The set of Nagios-based monitoring services necessary at the Resource Provider level and at the EGI central level is called the Service Availability Monitor (SAM). Tests to monitor all mission-critical infrastructure resources and services have to be defined and implemented as probes. A subset of probes will be able to raise alarms in the dashboard and are flagged accordingly. In the event of failure, notifications of the possible problem together with hints on how to solve the problem are sent to the technical staff and other relevant people allowing them to work on the problem before outages affect production and availability.

Alerts and warnings are delivered to the administrators via email and SMS, depending on the site managers' choice. Multi-user notification escalation capabilities ensure alerts reach the attention of the right people.

The execution of probes can be rescheduled to test the solution of a problem.

Statistical data is collected to provide input for the availability and reliability figures to see if OLAs are fulfilled and production level is reached. Only a subset of test results generating alarms in the dashboard is considered for the computation of monthly availability and reliability statistics.

A good monitoring system monitors not only the network and the resources, but also the accessibility and functionality of the used operational tools.

2.4.2 Requirements

1. Regionalization is an important factor since the Grid in its nature is a distributed system. Monitoring should therefore be split into various instances running in each region and a central instance collecting results. From the technical perspective the distributed system contributes to increased scalability as each instance covers a smaller number of Resource Centres than a single central instance. From the operational perspective, the Resource Providers get much more control and responsibility over the whole monitoring process since customization of the national monitoring infrastructure is under the local responsibility. This way, central problems no longer impinge local monitoring and response time should decrease by shortening the length of the reaction chain and removing a possible bottleneck. Finally, a distributed system enables individual instances to tune the monitoring by introducing extended custom probes to monitor custom services not covered by the generic profile. Also, individual instances can benefit from additional functionalities of the monitoring system such as direct email or text message notifications, extending monitoring on uncertified sites or direct scheduling of tests via a web interface.

2. Status and historical data should be accessible in a centralized portal. These historical records of outages, notifications, and alert response are relevant for later analysis.

3. The monitoring interface should also expose information for the calculation of resource availability and reliability.

4. Information shall be exchanged according to a given template and using a common transport mechanism (ActiveMQ).

5. It shall work as an input plug-in for the Operations Portal.

6. Additionally it would be desirable to not only monitor the resources but also the availability of the needed operational tools, such as the different regional monitoring instances.

2.4.3 Interoperability of different middleware stacks with SAM

The Service Availability Monitoring system is based on Nagios.

Nagios [R 38] is a well-known and mature general purpose monitoring system that enables organizations to identify IT infrastructure problems.

Out of the box, Nagios can already monitor many different infrastructure components - including applications, services, operating systems, network protocols, system metrics and network infrastructure. Furthermore, its extensible architecture allows easy integration with in-house and third-party applications. Hundreds of community-developed add-ons extend core functionality to ensure a faultless functioning of the entire infrastructure. New tests to monitor further mission-critical infrastructure components can be defined and deployed with freshly written probes for them.

Within EGI the central instance of SAM collects the monitoring results from the Resource Provider SAM instances, and provides a centralized MyEGI portal [R 39] to graphically display data, access status and historical data.

A dedicated central Nagios system ("ops-monitor") monitors the ActiveMQ Brokers network, the Resource Provider Nagios instances and other operational tools. CERN developed probes for monitoring these two services. The ops-monitor Nagios instance can be found at [R 40]. Additional probes for other operational tools are being developed.

To integrate a new middleware stack into Nagios, sensible tests for the service types defined in the management interface for this middleware have to be developed to cover the relevant functionality in the middleware stack. The probes are subsequently integrated into the SAM Release. For that the subset of probes which should raise alarms and have an influence on the reported availability and reliability metrics has to be defined. It may be sufficient to just have a compatible Nagios reporter from a different kind of monitoring tool which can be integrated in regional and central instances.

Since SAM Update-07 release (30th November 2010) SAM relies on the Aggregated Topology Provider (ATP). ATP is currently fed with information from both GOCDB and BDII. ATP extracts VO mappings from the BDII as those are not present in GOCDB. This is the reason why a top-level Information Discovery System which integrates different middleware stacks is paramount.

2.4.3.1 Tests and Nagios probes for gLite resources

Currently, the Nagios probes for the following gLite service types are available:

- APEL
- BDII (top and site BDII)
- Computing Element
- CREAM-CE
- FTS
- gRB/WMS
- Logging and Bookkeeping
- Local-LFC/Central-LFC
- MPI
- MyProxy
- SRM
- VO-box
- VOMS

Further documentation and descriptions on these probes can be found on [R 42].

Kommentar [Tiziana F1]: I don稚 understand this list. There are additional services that are monitored such as WMS for example. From which source of information was this list extracted?

2.4.3.2 Tests and Nagios probes for ARC resources

ARC probes are fully integrated with SAM [R 47] [R 49] starting from the release Update-7 and they monitor the ARC-CE service. The ARC GridFTP service will be monitored with probes for the standard GridFTP service. ARC-CE probes are maintained by the EMI ARC Product Team [R 52]. The ARC monitoring tests became operational on 7.04.2011.

2.4.3.3 Tests and Nagios probes for UNICORE resources

UNICORE probes are provided by NGI_PL [R 31]. The probes monitor the following services:
- Gateway
- UVOS
- Registry
- Workflow Service
- Service Orchestrator
- Global Storage

More details on those probes can be found in [R 53].
Integration of probes with the SAM is ongoing and planned for SAM Update-13.
Maintenance of UNICORE probes will be done by the respective UNICORE EMI product teams.

2.4.3.4 Tests and Nagios probes for Globus resources

Globus probes are fully integrated with SAM starting from SAM Update-12. The probes monitor the following services relevant for Globus:
- GSI-SSH
- GridFTP
- GRAM
- MyProxy.

Other services (e.g. LDAP) and basic checks (e.g. port checks and certificate lifetime) are covered by the same tests used for gLite services.
Maintenance of Globus probes is under the responsibility of the IGE project.

2.4.4 Procedures to integrate new Nagios Probes

There are some procedures in the Availability and Monitoring area. For the integration of new resources namely two of them are relevant:
- "Adding new probes to SAM" [R 54], approved by the OMB in March 2011, a procedure for adding new OPS Nagios probes to the SAM release.
- "Setting a Nagios test status to OPERATIONS" [R 55], approved by the OMB in November 2010: A Nagios probe is set to OPERATIONS when its results are used to generate notifications for the Operations Dashboard. This procedure details the steps to turn a Nagios test to OPERATIONS.

2.5 Accounting Interface

2.5.1 Functionality

The EGI Accounting Infrastructure 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 by dynamic queries on the parameters above at any level of the hierarchical tree structure which defines EGI and its partner grids.

Accounting is necessary to demonstrate that the usage of resources by user communities is in accordance with expectations. Site administrators are able to check actual usage of CPU resources against scheduling policies implemented at the site. VO resource managers are able to understand how CPU resources are utilized by their users.

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. This is assumed to be delivered by the exchange of accounting data at the appropriate level.

2.5.2 Requirements

An accounting interface has to fulfil the functionality described above. Further requirements are:

- Access to accounting data needs to respect all relevant policies and legal requirements. It is expected that this is controlled by the standard user authentication and authorization framework.
- Data identifying an individual should not be sent across the wide area network in plain text [R 45].
- As data from different grids is to be combined, a set of compatible units of measurement must be used. The CPU benchmarking tools currently in use are SPEC-INT 2000 [R 75] and HEP-SPEC 06 [R 46].

2.5.3 Current Status

The core EGI Accounting Infrastructure is based on APEL [R 34]. Other systems interface to APEL to collect data in one central place. The collected CPU accounting records are displayed in the Accounting Portal [R 43] as described above.

The bulk of existing Resource Centres collect data from their batch systems (e.g. LSF, Torque; SGE, Condor), which are joined with information about the job's user grid credentials and published to the central APEL repository. At the time of writing the EGI infrastructure is in transition of the transport layer from a private ActiveMQ broker to the production broker network already used by other EGI Operational Tools. The new system uses the STOMP interface to define a messaging model with encryption, verification, and acknowledgements. Other partner grids (Open Science Grid, IGI and NDGF), and a few additional Resource Centres with their own accounting services, currently publish summaries of data in the form described above directly into the APEL central repository. While participant Resource Infrastructures publish all of their VOs data, partner grids publish information for a subset of VOs (e.g. OSG).

CPU data is published in the form of either: job level records (JR) containing data from a single batch job; or summary aggregate job records (SJR) containing totals for a number of jobs run at a single Resource Centre for a single user and VO in a given month. The Job Usage Record (UR) schema is a plain text version of the OGF-UR v1.0 with some common extensions. For example, the original UR does not have the concept of a Resource Centre, which is on the other hand crucial. The summary record has been submitted to OGF's UR-WG for possible adoption as a community standard [R 35].

The OGF UR Working Group (UR-WG) is considering a proposal from EMI for a UR for storage accounting. 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 [R 56] to agree on the semantics of the existing UR and existing common extensions and possibly propose further extensions.

2.5.4 Integration with other Infrastructures

Other grid infrastructures who wish to publish accounting data need to:

a) Define a structure for their grid 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).

b) Extract data from their accounting system grouped 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.

c) Other infrastructures running a gLite CE (lcg-CE or CREAM) could run UMD software to aid collecting accounting data. Infrastructures running other middleware stacks who run one of the currently supported batch systems listed above can take UMD data collectors to parse the raw accounting data collected by the batch system to which they will then need to add the CPU speed and user/VO credentials, before publishing.

d) Register the publisher with APEL (by providing the host DN to the EGI APEL support unit). The APEL Repository only accepts accounting records from registered Resource Centres. For APEL client sites this is defined by the glite-APEL service type in GOCDB. An equivalent mechanism will be developed for summary publishing Resource Centres/Resource Infrastructures.

e) Publish the records into EGI's ActiveMQ Message Bus using the agreed encryption framework. The APEL repository will accept the records into a holding container from where they will be merged with the summaries from other grids 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.

f) From the master summary table, the data are then exported to CESGA where they can be viewed in the accounting portal.

2.5.4.1 Issues

- For the aggregation of user data it is assumed that all interoperating infrastructures use a user identity based on X.509 certificates signed by IGTF recognized Certificate Authorities.
- While a worldwide community management service like VOMS makes the aggregation of VO accounting data from different infrastructures simple, it would be feasible to implement a VO name transformation to combine the data from infrastructures who have named the same VO differently.
- Another issue is the unambiguous mapping of user accounts to VOs. In some cases users might belong to more than one VO in which case identifying to which VO the utilization results would go is not possible. Extra effort will be needed to check the fulfilment of arranged pledges.
- The issue of exchanging data identifying a user has been a contentious one. It is frequently asserted that this is illegal under the laws of certain countries. Extensive research was undertaken by the Joint Security Policy Group (JSPG) in EGEE-III during the development of the Grid Policy on the Handling of User-Level Job Accounting Data [R 45] with the result that legal advice was given that with the appropriate acceptable use policy and the agreement signed by the user and by the Resource Centre running the accounting repository, then the

collection, storage and restricted display of data identified by UserDN is acceptable. This issue might have to be re-evaluated again when exchanging accounting data with other infrastructures like e.g. DEISA [R 29].

- Current accounting is only of CPU of batch jobs but the interfaces between infrastructures should also allow the integration of other types of accounting record as they are developed.
- The currently agreed unit for normalization of CPU time in EGEE, EGI, and WLCG is HEPSPEC06 hours [R 46]. For interoperation with an infrastructure that does not collect this value from the resources running jobs, some conversion factor must be negotiated.

2.5.4.2 Future Work

At the time of writing the ActiveMQ interface into APEL only accepts a single type of job record for the CPU used by a batch job. The summary development mentioned above will include handling multiple types of record. As well as the summary record this will allow the repository easily to be extended to support other types of accounting, such as storage, as well as allowing evolution of the CPU UR. New accounting types should ideally be developed by all the infrastructures working together.

The RUS interface planned in APEL will allow other grid infrastructures to use a standard web services interface to publish records. This will replace item (e) in the integration list above.

For further discussion on accounting integration see [R 57]

2.5.4.3 ARC resources

Accounting integration was performed already during EGEE III. The aim was to gather and export accounting from the Nordic T1 and T2s, which for the compute part were based on ARC, and send data for selected VOs to the APEL central repository so they can be viewed with the EGI Accounting Portal. ARC-CE supports accounting via SGAS (SweGrid Accounting System, [R 19]) and an automatic script for exporting the accounting info gathered in SGAS to APEL was set up [R 20]. Currently, only LHC VOs are published to APEL but this could easily be extended to other international VOs.

The SGAS-APEL interface should be changed to the new once discussed above. This should be straightforward as the extraction and selection phase will not change, only the transport layer which will change from JDBC to ActiveMQ.

2.5.4.4 UNICORE resources

Currently no means of collecting accounting and usage records are directly implemented within UNICORE. Instead, this is done directly via the underlying batch system, see for example as in the DEISA project, where the accounting data is converted into OGF-UR format and provided according to XUUDB access control.

Accounting services for UNICORE have been developed by NGI_PL and NGI_BY. These are being reviewed within the UNICORE community. D-Grid within the NGI_DE is also building a regional service to collect accounting data from UNICORE and other clients. For all these implementations the common interface to publish data onwards to the EGI central repository needs to be used. Discussions have started with the developers on this.

2.5.4.5 Globus resources

IGE has adopted GridSAFE [R 58] as its accounting solution. It is currently under test. GridSAFE was designed as a site accounting repository to collect data locally but it has the interfaces to accept data from other Resource Centres too so it could act as a regional repository receiving data from a number of Resource Centres.

From the specification it does not have the ability to publish data on to higher levels in a hierarchy of repositories. It relies on others pulling data from it through an OGF RUS interface rather than the EGI push model. However a proof of concept was carried out in NGI_UK to use their Globus RUS client as a backend to GridSAFE to push data on to a remote RUS. This implies that data can be extracted so the APEL publishing model could be made to work.

2.6 Support Interface

2.6.1 Functionality

The user support infrastructure in use within EGI is 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 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 throughout the project enables to pass trouble tickets from one system to the other in a way that is transparent to the user. 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. A standard has been defined for the interface between ticket systems and also a template for a ticket layout exists to ensure the quality of service. These are documented in the GGUS documentation [R 36].

Ticket processing management (TPM) is responsible of ticket triage and holds a global overview of the state of all tickets. TPM is responsible for those tickets that have to be assigned manually, i.e. so that they get forwarded to the correct support units. TPM provides first-level support and keeps track of long-term trouble tickets and help to solve them with their very good general grid knowledge. In this way, a problem submitted to GGUS can be quickly identified as either a grid problem or a VO specific problem and addressed to the appropriate second line specialized support units or the dedicated VO support teams whose members have specific VO knowledge.

Second-level support is formed by many support units. Each support unit is formed from members who are specialists in various areas of grid middleware, or regional supporters for operations problems, or VO specific supporters. The membership of the support units is maintained on mailing lists.

2.6.2 Requirements

Regardless of the number of parties involved, the submitter of a trouble ticket should be able to transparently follow the chain of actions needed to solve the reported problem. This transparency together with the independence from the actual ticket system is used by the experts from the different areas who get assigned to the ticket. It can be seen that the main requirement of the ticketing system is that information flows between different parts of the EGI support network.

This is especially important since the support interface is not only used for 3rd level support dedicated to the end user, but also for the relevant parts of internal trouble ticket communication fulfilling standard operational, grid oversight and partially also development functionalities.

Other relevant requirements on the support interface is the existence of a functional body like the TPM as described above and the connection to a useful, searchable and well maintained knowledge base.

Other basic requirements that can be expected from a more advanced support ticket system:
- Differentiating between real problem tickets and service requests
- Ability to mark a ticket as spam

- Mail notification when a ticket is assigned to a support unit or person
- Possibility to involve several experts at the same time
- Searching tickets via ticket ID as well as via parameters
- Automatic reminders about open tickets
- Several tickets describing the same problem can be put into a master-slave relation.
- Other dependencies can be represented with child and parent relations.

2.6.3 Integration of new Resources into GGUS

There are three distinct cases to be considered when integrating new resources into the EGI user support infrastructure:

2.6.3.1 Integration of a new Resource Centre into the infrastructure

In case a new Resource Centre is added to the EGI infrastructure this is resources centre is always part of an NGI. For the user support area this is a simple case as the information about resource centres is extracted from GOCDB. This means that no manual steps are needed to integrate a new resource centre in GGUS.

2.6.3.2 Integration of a new NGI into the infrastructure

If a new NGI joins the EGI infrastructure it is required to provide a ticket system which is integrated with GGUS. This can be done in different ways, depending of the size and the maturity of the NGI.

- The simplest way, which might be suitable for a small new NGI is to use GGUS directly. This has the limitation of just one support unit for the whole NGI. Tickets cannot be assigned to specialized groups or specific resource centres within the NGI. This further processing of the tickets is done independently from the EGI support infrastructure.
- The NGI can make use of xGUS, which is a customisable slimmed-down regional instance of GGUS. xGUS is hosted and maintained by the GGUS team. Customization can be done via an administrative web interface, which enables creating and managing support units and defining special workflows. xGUS comes with the interface to GGUS built in.
- The NGI can set up an own ticket system. In this case the NGI has to make sure that their ticket system fulfils the requirements of the interface definition to GGUS. The NGI ticket system needs to be interfaced to GGUS and the NGI is responsible for maintaining this interface.
- Details on the NGI creation process are documented in a specific procedure [R 37].

2.6.3.3 Integration of a new Technology Provider into the infrastructure

Should EGI decide to utilize software from a technology provider that has not so far involved with the project, an agreement has to be found with that technology provider on how to integrate its support infrastructure within the EGI Helpdesk. This process is already complete for the EMI and IGE projects. EGI has set up a Technology Helpdesk which is interfaced to GGUS for that purpose. No general description of the details of the integration of a new technology provider into the Technology Helpdesk can be given here, as this is highly dependent on the internal support structure of the respective technology provider. Nevertheless it is important that this is done in a way that enables EGI to have an overview of issues with the products provided by the technology provider and to gather statistics on the quality of the support given by the provider.

EMI has set up a structure within the Technology Helpdesk for its various products, including ARC, gLite and UNICORE.

3rd level support for Globus will be provided by IGE. IGE provides a support infrastructure for the European Globus users in all European, national, and regional e-Infrastructures with EGI and DEISA/PRACE being the most important ones. The Technology Helpdesk contains a queue to forward 3rd level support tickets directly to the IGE user support team.

For details on the Technology Helpdesk refer to [R 59].

2.7   Dashboard Interface

2.7.1   Functionality

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 combines and harmonizes different static and dynamic information and therewith enables the operators to react on alarms, 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.

2.7.2   Requirements

A dashboard interface has to fulfil the functionality described above.

With the increasing relevance of the SAGA Service Discovery specification (OGF) [R 25] for a standards-based approach for interoperability one more requirement on the dashboard is to provide such a well defined interface in order to be prepared for the harmonized integration of many different third party information providers.

We assume that EGI as a whole should try to unify the input from Resource Centres, which should publish their information via a harmonized and unified Information Discovery System based on GLUE 2.0 and in a generalized form of BDII. In addition, access should be limited to users that are authenticated through a common user authentication system such as VOMS (see also section 2.8).

2.7.3   The Operations Portal

The Operations Portal [R 23] content is based on information which is retrieved from several different distributed static and dynamic sources – databases, the EGI Information Discovery System, web services, etc. – and gathered onto 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 implementation that provides a transparent integration of each of these resources. The web service in question is named Lavoisier [R 24].

The goals of Lavoisier are to provide:

* a web layer as independent as possible from the mechanisms technology used to retrieve the original information,
* intermediate information usable in the same format in order to cross-query,
* information which is independent from the availability of the data provider.

This solution design means that the web application does not need to know the exact location of the data provider and neither which kind of technology has provided the information initially. All these concerns are already taken into account by Lavoisier.

Lavoisier has been developed in order to reduce the complexity induced by the various technologies, protocols and data formats used by its data sources. It is an extensible service for providing a unified view of data collected from multiple heterogeneous data sources. It enables us to easily and efficiently execute cross data sources queries, independently of used technologies. Data views are represented as XML documents and the query language is XSL.

The global architecture of the Operations Portal is based on a plug-in schema, where information can be retrieved from heterogeneous data providers. These plug-ins transform information in various formats extracted from different technologies (i.e. RDMS, JSON, JMS, ldap, http, web service) into a standard format XML. At this stage it is easy to execute cross data sources queries by using XSLT transformation. In the end the web application is using all information in the same format (XML).

2.7.3.1 Integration of a new resource

The architecture of the portal has been designed to propose a standard access to information from an extended number of data sources. The integration of new data sources is eased by the use of the Lavoisier web service.

In the case of a known technology we will create and add a new view by using an existing plug-in out of the wide-range of plug-ins already available.

If a site and its resources are already integrated in all the other operational tools through existing information providers (e.g. registered in GOCDB, monitored by Nagios, publishing their information via BDII and having a tree in GGUS), existing plug-ins can be reused and no additional integration effort for the usage of the Operations Portal is needed. For new providers, new plug-ins can be developed as needed.

The integration of different information systems present in different middlewares such as ARC, UNICORE, or Globus can be done via an abstraction layer.

One such a possible abstraction layer could be to integrate the SAGA Service Discovery specification (OGF) [R 25] into a Lavoisier plug-in which will permit to access information using different services (like the information service of UNICORE – CIS [R 26]) and different schemas like CIM [R 27] or the GLUE standard [R 14].

The modularity of Lavoisier allows the easy integration of almost any kind of information. Such integration is certainly needed and meaningful for the new resource types entering EGI, such as HPC systems, virtualized resources or desktop resources. As long as these resources are monitored, it is possible to integrate them via plug-ins inside Lavoisier. The integration will be done step-by-step during the whole project according to the identified priorities.

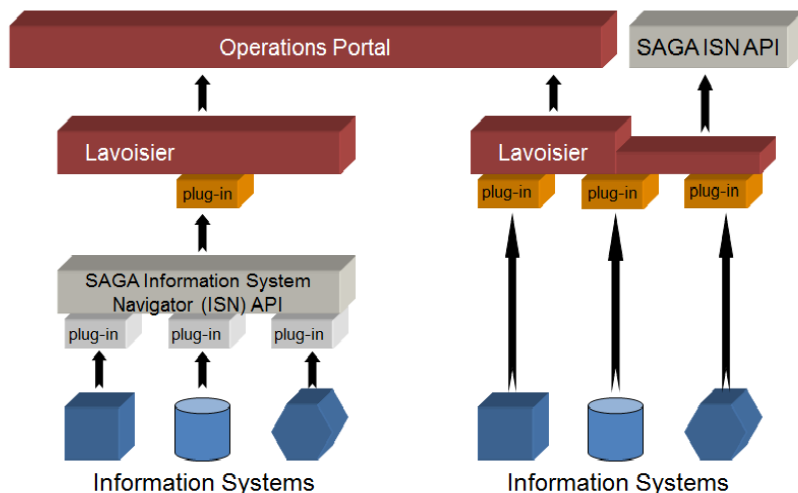### 2.7.3.2 Alternative possibilities to integrate new information providers

Figure 1. Integration of new information systems into the Operations Portal

The alternative depicted on the left side of Figure 1 might seem more work at first, but part of this can be outsourced to the information providers and reused for other purposes. On the other hand, a Lavoisier to SAGA Information System Navigator (ISN) link might be needed anyway. The two implementation solutions can coexist and might be combined.

2.7.3.3 Integration of gLite resources
Plug-ins for all relevant information providers in the case of a site's gLite resources (Nagios, GOCDB, GGUS, BDII) exist and gLite resources can therefore be operated from within the Operations Portal.

2.7.3.4 Integration of ARC resources
Plug-ins for all relevant information providers in the case of a site's ARC resources (Nagios, GOCDB, GGUS, BDII) exist and gLite resources can therefore be operated from within the Operations Portal.

2.7.3.5 Integration of a UNICORE resources
The UNICORE resources are registered in GOCDB and the integration with SAM is in progress; the GGUS trees exist. Hardware information following the GLUE standard could be taken from the Central Information Service CIS over the SAGA ISN API link.

2.7.3.6 Integration of a Globus resources
Globus GT5 resources are registered in GOCDB and the integration with SAM is in progress; the GGUS trees exist.

Taking into account that LCG-CE is very similar to Globus GRAM, lcg-ce information providers can be reused for the BDII. With that Globus resources should be able to be directly integrated into the operational dashboard. The issue of the integration of GLOBUS into a unified Information Discovery System was discussed at the TCB, and is being investigated.

2.8 User Membership Management, Authentication and Authorization

The actual way user authentication and VO membership management effect many operational interfaces that have been defined so far. This might be especially true for accounting, but is equally relevant for monitoring or when using a high level tool like the operational portal.

The basic information on who is authorized to access resources and services operated in a Resource Centre can be stored in different ways within different distributed infrastructures interested to join or collaborate with EGI.

Within the EGI production infrastructure X.509 certificates and its proxy derivatives are used for user authentication. A user would e.g. request an X509 credential with VOMS extensions from a national or organizational Certificate Authority (CA) which is recognized by the International Grid Trust Federation (IGTF) (see also [R 11]). Resources within the production infrastructure are made available to users depending on their VO membership. Access to such a VO is governed by a VO Manager who is responsible for managing the addition and removal of users and the assignment of users to groups and roles within the VO.

Normally in a VO, the VO Manager has the authority to manage user membership and roles. In order to control access in a more fine grained manner (for example to ban users, or limit the access to some of the resources) an authorization service is needed (ARGUS) which holds information on how to map users to local accounts.

In EGI there are resource providers who are not willing to offer pool accounts on their resources in order to enforce proper access control. Users have to apply for a personal account first and have a certificate mapped to it.

However, there are alternative ways to distribute authorization information across a grid infrastructure. In D-Grid for example a centralized approach is used: the central Grid Resource Registration Service (GRRS) knows about resources and which VOs are allowed to use them. Each VO has a VO Management Registration Service (VOMRS) server where users are registered with their certificate and D-Grid userID after they have applied for a userID and the VO membership. From this information a service prepares mapping files for Globus, gLite, dCache [R 7], and UNICORE for each Resource Centre. Such files used by the relevant local services, e.g. the UNICORE User Database XUUDB. Alternatively, for UNICORE Resource Centres information can be maintained in the UVOS service[1].

Within EGI the harmonization of user authentication and authorization will rely on the work plan of EMI.

2.8.1 Desired Functionality of a user authorization system

- Providing a consistent approach for identical DN/UID mapping.

- Global banning and unbanning of users over sites and services

- Providing an administrative tool to maintain and control DNs and policies, especially also supporting hierarchical policies.

2.8.2 Requirements on a user authorization system

Basic requirements that can be required from a user authorization system and which are relevant for the integration:

- Identical user mapping functionality

---

[1]   The usage of the UVOS service is the solution of choice of PL-Grid.

Kommentar [Krzysztof2]: This very centralized approach is one possibility but sometimes it is not flexible enough. This is a contrary solution to ARGUS EES, where there is a separate service with information how to map users to local accounts. BTW in UNICORE such information (for all sites, for specific sites, for groups of sites) can be more easily kept in UVOS: in PL-Grid this is used. AFAIU D-Grid wanted to have one simple solution for different stacks.

Kommentar [Michaela 3]: From Krzysztof Benedyczak:
Whole intro referring to VOMS ACs only. For UNICORE it won't be used - SAML assertions (possibly issued by VOMS) carry attributes. This is especially contradictory to the 2.8.3 section when the document quietly shifts to SAML assertions. Yep - in case of gLite either middleware must be converted to SAML (personally I doubt it will happen) or translation of credentials must be made.

- It should be possible to use a centralized approach to do the DN/UID mapping in a consistent manner which provides a good level of abstraction for users so they will not be involved in dealing with low level details of platforms.

- Policy based user access
  - Site administrators should be able to ban users based on DNs, CAs, VOs for the whole site or over multiple services.
  - The banning list and other policies can be created and written down in a well defined way, e.g. by using a language to create and customize policies.

- Support for single-user and multi-user pilot jobs
  - Pilot jobs are submitted through pilot submitter and the real owner of the jobs until they start execution on the worker nodes are not known which is important in the case of accounting. Using a service, it should be possible to map users to a particular POSIX UID/GID.
  - This requirement is possibly not equally urgent as the other two, since authorization problems are only expected for multi-user pilot jobs.

2.8.3 ARGUS

EMI has selected the ARGUS authorization framework as general approach for user authorization based on the common SAML profile which shall be supported over all middleware stacks.

ARGUS is an authorization system for distributed services such Compute Elements, Portals and Worker Nodes and it replaces the Site Central Authorization Service (SCAS) as used in different gLite tools and several non-Webservice based Globus Toolkit 4 components. In order to achieve this consistency a number of points must be addressed. ARGUS consists in several distinct components (Figure 2). The first component is the Policy Administration Point (PAP for short) service where all policies are defined and stored. Second, authored policies must be evaluated in a consistent manner; this task is performed by the Policy Decision Point (PDP). And finally, the data provided for evaluation against policies must be consistent; this is done by the Policy Enforcement Point (PEP). The interfaces to the PAP and PDP daemons are standardized and well defined. PEP uses a proprietary protocol.

The eXtensible Access Control Markup Language (XACML) [R 62] is a declarative access control policy language based on XML and can be used as a processing model which describes how to interpret the policies. The EMI XACML working group is aiming at standardizing the XACML attributes [R 61] used in the requests.
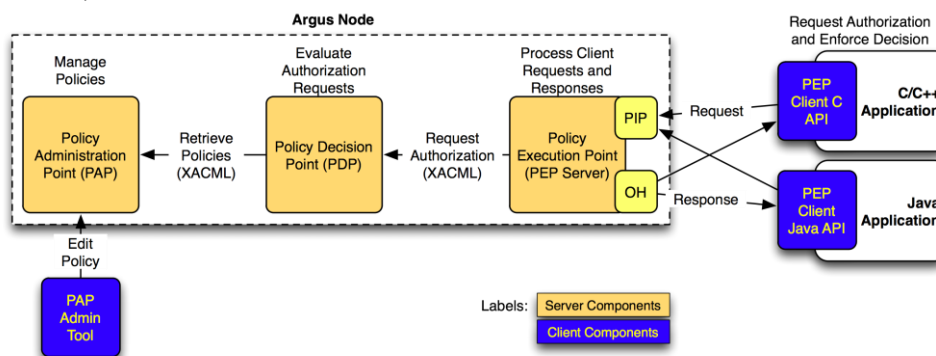
Figure 2: Internal ARGUS Components

The three, so far presented, ARGUS components (PAP, PDP, PEP) are responsible for authorization. ARGUS-EES is the component which maps DNs to particular POSIX UID/GID. It is normally contacted by the PEP. But not all middleware stacks are using PEP. It has also to be noted that ARGUS supports hierarchical policies since a PAP can use another PAP.

2.8.3.1 ARGUS and gLite

Several services can interact with ARGUS in gLite; eventually every service that uses SCAS for users' validation can be migrated to use ARGUS. The site policies are maintained using the command pap-admin. By default Argus contains an empty policy and no one will be permitted to do anything. Basically ARGUS is designed to answer questions in the form of *Can user X perform action Y on resource Z at this time?.* If so, ARGUS gives a response to the PEP java client and the user can perform the action. If the request does not match to any appropriate access control policy then the access is rejected. Each policy is evaluated from most to least recent, the first policy that matches is the result returned by Argus. As example, if the first policy is a policy that would deny the access and then a new one is added that would permit it, the result of an authorization request will be permit since the permit policy is most recent.

Several gLite services are/will be integrated with the ARGUS EMI authorization system:

- CREAM: ARGUS policies grant access to grid users to access CREAM-CE computing resources. When a new user job is submitted to CREAM the site ARGUS instance is requested to accept or deny the job submission based on the site ARGUS policy.
- WN/gLExec: Pilot jobs can be mapped to a specific grid user based on ARGUS policy response instead of SCAS. Pilot jobs are mapped to grid users into WN using the LCMAPS C PEP Plug-in to contact the Argus framework. In the Argus deployment scenario (similar to the SCAS deployment scenario) the LCAS framework is redundant. In future releases of gLExec the LCAS framework can be switched off and in a later stage complete be removed from the system.
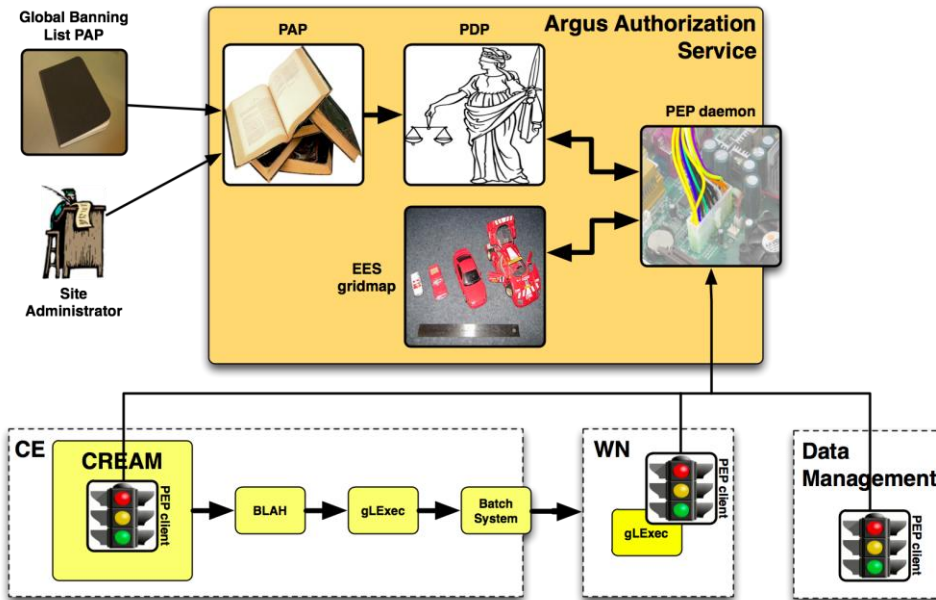
Figure 3: gLExec and Argus integration

2.8.3.2 ARGUS and ARC

ARC middleware requires a consistent mechanism to provide authorization based on user DNs. Existing ARC releases don't provide coherent solutions to address issues such as identical DN/UID mapping, DNs and policy maintenance, Global banning and unbanning of users over sites or specific services and support for accounting of pilot jobs. To overcome these issues, the ARGUS authorization framework is integrated within the Hosting environment daemon (HED) component in ARCv1.

HED is in charge of authorization requests for incoming user jobs. During the user ID mapping process the HED component initiates the authorization client which then communicates with the PEP daemon in ARGUS. As a first step, the ID mapper within HED collects the Grid credentials and tries to configure the HED authorization client so it can establish a communication channel between the HED client and the ARGUS authorization framework to send and receive the XACML requests/responses.

By default an ARC authorization and authentication request is composed of a XACML subject, resource, action and an additional XACML environment element which differs from the response structure received by ARGUS with attributes such as: XACML decision element and obligation. The HED authorization client uses the gLExec LCMAPS plug-in to send and receive these requests and eventually parse the XACML response decision to authorize the user and the obligations to map a user to a local account.

Currently as a proof of concept an ARGUS client is in charge of sending/receiving messages to the PEP daemon. However, eliminating communication to the PEP daemon from the ARC authorization client will increase the performance and can be achieved through providing a profile for ARC in ARGUS.

Further details on implementation, deployment and configuration examples can be found in [R 63].

2.8.3.3 ARGUS and UNICORE

For the case of UNICORE what normally is referred to as authorization is split into two terms: "authorization" in UNICORE means the decision if a certain request is allowed or not; "incarnation" in UNICORE means to map a request to a local system (what includes more than in e.g. in the case of gLite: not only UID/GID(s), but also symbolic application names are mapped, as well as symbolic arguments, execution environments, etc.).
UNICORE has a built in mechanism called PDP which is responsible for the actual authorization. The administrator can choose its implementation. The default implementation uses a file based authorization policy. This default XACML based policy predefines attributes to allow/ban a user. Therefore authorization is typically administrated by assigning attributes for users using tools of choice: UVOS, XUUDB, files. XACML policy is modified only in case of complicated use-cases (e.g. banning all users of a certain VO but only at night). So in the case of UNICORE authorization can already be controlled to the desired level without using ARGUS. ARGUS can be seen as an intermediate solution: its usage will allow for more flexibility than is provided by assigning attributes while still allowing administrators not to learn a complicated XACML syntax. However a really advanced authorization problem will still require manual XACML policy editing. ARGUS integration may also be considered if grid deployments (because of e.g. legacy reasons) prefer to keep attribute sources very simple.
As to incarnation, attribute source services (UVOS/XUUDB/or even a file) define permitted and default values for users/groups of users etc. within UNICORE. As in the case of e.g. D-Grid the input and definition files for these attribute source services can be created in a more global way. Additionally a local configuration file is used for application related data. Users can express preferences to choose desired values (e.g. a desired GID) out of possible ones. Additionally the local administrator can define hooks which modify the incarnation.

So even if the current user management already fulfils our basic requirements it will be useful to integrate UNICORE with an EGI wide supported user authorization system for the sake of unified access or in scenarios where different middlewares are deployed on one Resource Centre.

In order to integrate ARGUS with UNICORE there are three different integration options to be discussed:

1. Usage of ARGUS PDP: As with version 6.4.0 UNICORE can be configured not to use the local XACML file as in the default implementation, but to contact ARGUS PDP instead. The ARGUS PEP component can be skipped. The drawback of this approach is that a web service is needed for each request at the cost of some performance penalty.

2. The ARGUS PAP is used directly. A prototype is being developed by EMI and will be part of version 6.4.2. Policies are fetched from the ARGUS PAP and evaluated locally. This solution is tolerant to a failure of the ARGUS server.

3. A third integration concerns the use of ARGUS EES for incarnation. To do so a refactoring of the UNICORE container is needed. This feature is being planned; currently only the UNICORE native incarnation is possible.

2.8.3.4 ARGUS and Globus

Globus has a legacy of being the de facto toolkit to build upon and construct grid middleware clients and services. The infrastructures that use Globus without modifications or add-ons will only be able to authorize their users using a grid-mapfile. Depending on their infrastructure setup, the Globus services can authorize their users based on either a Unix account, a Kerberos account or an X.509 certificate, using a grid-mapfile.

The Globus services Gatekeeper, GridFTPd and GSI-OpenSSHd are well known for their support of the grid-mapfile and are still used as core-services in many Grid infrastructures. Many grid infrastructures have extended Globus by introducing features like the pool account support, i.e. assigning non-personal Unix accounts to users based on their credentials, and the support to authorize and map accounts based on VOMS credentials.

The extensions build upon Globus are being adopted as supported integrations through the IGE project. In effect the native Globus infrastructures will gain the ability to use pool accounts and VOMS based authorization as LCAS and LCMAPS implement these features. The LCAS and LCMAPS framework can extend the services with a pluggable (security focused) framework that can be extended with third-party plug-ins. In terms of feature implementations the Globus infrastructures will be on par with the CREAM CE and LCG-CE compute services.

On the roadmap is the integration with Argus to extend the authorization capabilities of LCAS and LCMAPS. In a similar way as the previous extensions, the Argus framework will complement and extend the authorization capabilities with the richness of the XACML policy engine and infrastructure potential of Argus to connect multiple PAP services together between sites.

The integration of Argus in Globus will be very similar to existing Argus integrations seen in other middleware stacks. Access to the compute facility or storage will be authorized by the Argus service based on the active XACML policy. When a user accesses a compute or storage facility, her credentials will be used in the authorization request to an Argus service node. The response will contain the authorization decision and the Unix account to which the credentials are mapped to. The adoption of Argus in the core-services will be transparent to the users.

The Globus strategy is to offer complete software solutions, starting with the core-services. Other service from Globus will gain a similar integration with Argus when the core-services are released with its support.

3 PROCEDURES AND POLICIES

Compliance to procedures and policies is important to ensure seamless interoperation of operations across EGI. These are needed to guarantee that OLAs are fulfilled. OLAs are a precondition for a high quality and stable production environment.

Procedures need to be independent from any actual operational tool used, and have to be middleware-agnostic. EGI procedures can be complemented by extensions that are specific to the needs of the Resource Providers.

Similarly, EGI security policies are formulated in general terms in order to be adopted by different infrastructures. Different infrastructure providers like e.g. DEISA adopted them by complementing them with several add-ons. An example of common security policy is the Acceptable Use Policy (AUP).

The Infrastructure Policy Group (IPG) regularly updates these documents and ensures communication between the different partners.

3.1 Current EGI Procedures and Policies

The operational procedures used within EGI have been evolved and further developed from those established within the EGEE series of projects and now have broad community support and adoption. Correspondingly, all current procedures and related operational work flows are directly reflected within the Operations Portal. As a result, the portal has to be regularly updated as the procedures change to reflect the needs of the community.

The EGI procedures and policies are collected in the EGI wiki [R 10]. These are approved by the OMB and periodically reviewed.

Procedures and policies are complemented by manuals, best-practices and how-TOs [R 64] and [R 65].

One procedure explicitly worth mentioning, since it has a great impact on the integration of new resources into the monitoring interface and the quality assurance of those new production resources, is the procedure for turning a SAM test into OPERATIONS. This procedure defines which tests are able to generate a notification in the dashboard in case of error and which are used to calculate the availability league table.

EGI has three security-related procedures:

- EGI Security Incident Handling Procedure [R 76],
- EGI Software Vulnerability Issue Handling Procedure [R 77],
- EGI-CSIRT Critical Vulnerability Operational Procedure [R 78].

In the deployed EGI infrastructure all problems concerning security should be dealt with between the EGI Computer Security Incident Response Team (CSIRT) and the EGI Software Vulnerability Group (SVG), [R 41].

CSIRT advises the resource centres on security matters and has the power to suspend them from the infrastructure if they fail to apply critical security patches.

The EGI Incident Response Task Force (IRTF) makes sure that incidents are handled according to the Incident Response Procedure.

The SVG ensures that the software available for installation on the EGI infrastructure is sufficiently secure and contains as few vulnerabilities as possible, thus reducing the likelihood of incidents.

When introducing a new technology in UMD one representative of the new Technology Provider has to be appointed to participate to SVG and to the Risk Assessment Team (RAT).

RAT is the group of people within SVG who carry out the issue handling process of the SVG, and are party to information on vulnerabilities which are not disclosed publicly. The RAT members are

developers from the various Technology Providers whose software is part of UMD, representatives of NGIs and experienced Resource Centre administrators.

4 FUTURE PLANS

The functionality and the requirements of the different operational tool interfaces described in this milestone will evolve over time. Operational requirements will continue to be collected from Resource Providers that are interested in integrating novel resource types into their e-Infrastructure as required. Input from infrastructure providers planning to operate different middleware stacks will also be gathered. In parallel to this, the integration with other Distributed Computing Infrastructures will likely bring new requirements for the extension of the operational interfaces currently deployed in EGI for monitoring, accounting, communication, management and support, as well.

The second year of the project will be focused on the completion of the integration of UNICORE and Globus, and on the integration with desktop Grids and PRACE.

As to the integration with desktop Grids, various possibilities are being investigated. In particular, the desktop Grids are being consolidated as operationally unified infrastructure, and the signing of a Resource Infrastructure Provider MoU with EGI is being discussed.

The integration with PRACE is being driven by user communities that require the coupling of high throughput and high performance computing. A pilot is being implemented in collaboration with the MAPPER project which comprises a selected list of EGI Resource Centres and PRACE centres. A joint EGI/PRACE task force [R 79] was constituted to foster progress of this integration activity.

The provisioning of virtualized services is being discussed with the user community, Resource Providers and Resource Centres. An EGI workshop was organized in May 2011, and use cases will be further discussed at the EGI Technical Forum in Lyon in September 2011. A task force was constituted in August 2011 to steer the discussion of use cases, implementation aspects and the operational integration of virtualized resources.

Kommentar [Michaela 5]: *Reply to Tiziana Ferrari (2011-08-18, 17:34): "..."* I just wonder where the UNICORE and Globus task forces are now mentioned. Didn't you want to have them mentioned more explicitly, Tiziana? At least they are a good example on how the whole integration process can be done efficiently.

5 REFERENCES

R 1	MS405: Operational Security procedures https://documents.egi.eu/secure/ShowDocument?docid=47
R 2	JRA1 WP7: Operational Tools Description of Work summary https://wiki.egi.eu/wiki/WP7:_Operational_Tools_DoW_summary
R 3	Integration of EMI support units into GGUS https://twiki.cern.ch/twiki/bin/view/EMI/MilestoneMSA11
R 4	StratusLab http://stratuslab.eu
R 5	Joint Security Policy Group, JSPG http://www.jspg.org/
R 6	MS407 Integrating Resources into the EGI Production Infrastructure https://documents.egi.eu/document/111
R 7	dCache http://www.dcache.org/
R 8	LFC catalogue service http://goc.grid.sinica.edu.tw/gocwiki/How_to_set_up_an_LFC_service
R 9	VOMS https://twiki.cnaf.infn.it/twiki/bin/view/VOMS/WebHome
R 10	EGI Operational Procedures https://documents.egi.eu/secure/ShowDocument?docid=209 https://wiki.egi.eu/wiki/Operational_Procedures
R 11	EGI Trust Anchor distribution https://wiki.egi.eu/wiki/EGI_IGTF_Release_Process
R 12	M.Ellert et al., Future Generation Computer Systems 23 (2007) 219-240.
R 13	Field L and Schultz M W Proc. of CHEP 2004, CERN-2005-002, 2005
R 14	GLUE Schema Specification http://www.ogf.org/documents/GFD.147.pdf
R 15	gLite WMS http://glite.web.cern.ch/glite/packages/R3.1/deployment/glite-WMS/glite-WMS.asp
R 16	European Desktop Grind Initiative (EDGI) Project http://edgi-project.eu/
R 17	Software Provider SLA Agreement https://documents.egi.eu/secure/ShowDocument?docid=212
R 18	UNICORE bug tracker http://sourceforge.net/tracker/?group_id=102081&atid=633902 UNICORE feature tracker http://sourceforge.net/tracker/?group_id=102081&atid=633905
R 19	SweGrid Accounting System SGAS http://www.sgas.se
R 20	SGAS to APEL Byrom R et al. http://www.gridpp.ac.uk/abstracts/allhands2005/apel.pdf
R 21	EMI software maintenance and support plan https://twiki.cern.ch/twiki/bin/view/EMI/DeliverableDSA11

R 22	Towards sustainability: An interoperability outline for a Regional ARC based infrastructure in the WLCG and EGEE infrastructures, L Field et al., 2010 J. Phys.: Conf. Ser. 219 062051
R 23	Operations Portal New Home Page http://operations-portal.egi.eu/
R 24	Lavoisier Home page http://grid.in2p3.fr/lavoisier
R 25	SAGA Service Discovery API http://www.ggf.org/documents/GFD.144.pdf
R 26	Common Information Service (CIS) for UNICORE Grids http://www.unicore.eu/community/development/CIS/cis.php http://www.d-grid.de/fileadmin/user_upload/documents/MonitoringWorkshop/Memon.pdf
R 27	Common Information Model Home Page http://www.dmtf.org/standards/cim/
R 28	UNICORE support mailing lists for EMI related and general issues: emi-support@unicore.eu and unicore-support@lists.sourceforge.net.
R 29	DEISA user management http://www.deisa.eu/services/user-related#usermngt
R 30	UNICORE 6 Monitoring with Nagios http://www.d-grid.de/fileadmin/user_upload/documents/MonitoringWorkshop/Rambadt.pdf
R 31	PL-Grid UNICORE Monitoring System http://www.unicore.eu/summit/2010/presentations/18_Bala_Monitoring.pdf
R 32	UNICORE architecture http://www.unicore.eu/unicore/architecture.php
R 33	Relational Grid Monitoring Architecture http://www.r-gma.org/
R 34	APEL Home https://wiki.egi.eu/wiki/APEL
R 35	Extensions to OGF-UR V1.0 as used in APEL http://forge.ggf.org/sf/docman/do/listDocuments/projects.ur-wg/docman.root.current_drafts.aggregate_ur_schema
R 36	GGUS Documentation on interfaces https://ggus.eu/pages/ggus-docs/interfaces/docu_ggus_interfaces.php
R 37	NGI Creation Process https://wiki.egi.eu/wiki/PROC02
R 38	Nagios http://www.nagios.org/documentation
R 39	MyEGI Portal https://grid-monitoring.egi.eu/myegi
R 40	Ops-monitor Nagios instance https://ops-monitor.cern.ch/nagios
R 41	Security Policy Group https://wiki.egi.eu/wiki/SPG Software Vulnerability Group https://wiki.egi.eu/wiki/SVG
R 42	Nagios Probe Documentation and Description https://tomtools.cern.ch/confluence/display/SAM/Probes https://wiki.egi.eu/wiki/Operations:Operations_tests
R 43	EGI Accounting portal http://accounting.egi.eu/
R 44	WS J. Ainsworth, S. Newhouse, and J. MacLaren. Resource Usage Service (RUS) based on WS-I Basic Profile 1.0. UR, August 2005

R 45	Grid Policy on the Handling of User-Level Job Accounting Data https://edms.cern.ch/document/855382
R 46	HEP-SPEC06 https://hepix.caspur.it/benchmarks/doku.php http://hepix.caspur.it/afs/hepix.org/project/ptrack/#SPEC_CPU2006
R 47	Integrating ARC into SAM https://tomtools.cern.ch/jira/browse/SAM-751
R 48	Building Packages on the SA1 Koji build system https://twiki.cern.ch/twiki/bin/view/EGEE/EGEESA1BuildingPackages
R 49	Additional steps required when supporting ARC services https://tomtools.cern.ch/confluence/display/SAM/Setup+SAM+for+ARC+services
R 50	GOCDB Documentation Index https://wiki.egi.eu/wiki/GOCDB/Documentation_Index
R 51	GOCDB/Input System User Documentation: Service Types https://wiki.egi.eu/wiki/GOCDB/Input_System_User_Documentation#Service_types
R 52	Nagios Tests http://wiki.nordugrid.org/index.php/Nagios_Tests
R 53	UNICORE Monitoring Infrastructure Probes http://alfred.studmat.umk.pl/~szczeles/PL-Grid/UMI-Probes.html
R 54	Adding new probes to SAM https://wiki.egi.eu/wiki/PROC07
R 55	Procedure for setting Nagios test status to operations https://wiki.egi.eu/wiki/PROC06
R 56	EMI compute accounting record https://twiki.cern.ch/twiki/bin/view/EMI/ComputeAccounting
R 57	Operational Tools Accounting Work Plan https://documents.egi.eu/public/ShowDocument?docid=531
R 58	GridSAFE project http://www.jisc.ac.uk/whatwedo/programmes/eresearch/gridsafe.aspx
R 59	EGI Helpdesk and the NGI Support Units https://documents.egi.eu/public/ShowDocument?docid=522
R 60	EMI SAGA Service Discovery http://hepunx.rl.ac.uk/edg/sa3-uk/sd/
R 61	Common XACML Authorization Profiles https://twiki.cern.ch/twiki/bin/view/EMI/EmiJra1T4XACML
R 62	eXtensible Access Control Markup Language, (XACML) Version 2.0, http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-core-spec-os.pdf
R 63	ARC integration with Argus http://wiki.nordugrid.org/index.php/Argus_integration
R 64	EGI Operations Manuals https://wiki.egi.eu/wiki/Operations_Manuals
R 65	EGI Best Practices https://wiki.egi.eu/wiki/Operations_Best_Practices
R 66	Security Incident Response Policy https://documents.egi.eu/public/ShowDocument?docid=82
R 67	Resource Centre Registration and Certification Procedure https://wiki.egi.eu/wiki/PROC09
R 68	Availability Re-computation Policy https://tomtools.cern.ch/confluence/display/SAM/Availability+Re-computation+Policy

R 69	UMD Release Schedule https://documents.egi.eu/public/ShowDocument?docid=526
R 70	Group unicore-integration-tf https://www.egi.eu/sso/groupView/unicore-integration-tf
R 71	Uniform Resource Identifier (URI): Generic Syntax http://www.ietf.org/rfc/rfc3986.txt
R 72	Group globus-integration-tf https://www.egi.eu/sso/groupView/globus-integration-tf
R 73	UNICORE Integration https://www.egi.eu/indico/categoryDisplay.py?categId=49
R 74	GLOBUS Integration https://www.egi.eu/indico/categoryDisplay.py?categId=53
R 75	Standard Performance Evaluation Corporation http://www.spec.org/
R 76	EGI Security Incident Handling Procedure https://documents.egi.eu/document/710
R 77	EGI Software Vulnerability Issue Handling Procedure https://documents.egi.eu/document/717
R 78	EGI-CSIRT Critical Vulnerability Operational Procedure https://documents.egi.eu/document/283
R 79	MAPPER-PRACE-EGI Task Force (MTF) https://wiki.egi.eu/wiki/MAPPER-PRACE-EGI_Task_Force_(MTF)