



EGI-InSPIRE

TECHNICAL NOTE

Observations and assessment of user community requirements from MS305

Document identifier:	TechNote on MS503 v3
Date:	03/12/2010
Activity:	NA3, SA2
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/267

Abstract

The requirements collected by NA3 from the EGI User Community and other sources are captured in MS305. This document provides an analysis of these requirements by categorising them with respect to the EGI Capabilities described in the UMD Roadmap and includes prioritisation information from the recent UCB meeting on 30th November 2010. This information will be used to inform the middleware prioritisation discussions with the Technology Providers that will take place at the TCB on 6th December 2010.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Michel Drescher, EGI.eu (MD) Gergely Sipos, EGI.eu (GS)		
Reviewed by	Steven Newhouse, EGI.eu Steve Brewer, EGI.eu		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
0.5	1 Dec 2010	First draft	MD / EGI.eu
1	2 Dec 2010	First draft including comments	MD / EGI.eu
1.1	2 Dec 2010	Suggestion for changes Merging UCB section	GS / EGI.eu
1.2	2 Dec 2010	Review and improving text	SN / EGI.eu
2	3 Dec 2010	Reviewing comments, page references to requirements	MD / EGI.eu
3	3 Dec 2010	Formatting, cross-checking with MS305_v6	GS / EGI.eu

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:
<https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY

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



VII. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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

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



VIII. EXECUTIVE SUMMARY

This document provides a brief summary analysis of the requirements collected and presented in MS305 User Feedback and Recommendations [R 1]. It categorises the relevant technical requirements according to the EGI Capabilities defined in the UMD Roadmap [R 2] leaving out those requirements that relate purely to User Community services or other aspects of the infrastructure. The recurring Technology related requirements are presented alongside prioritisations that have been endorsed by the EGI User Community Board on the 30th November 2010 and collected here for input into the TCB discussions on 6th December 2010.



TABLE OF CONTENTS

1	INTRODUCTION	6
2	REQUIREMENTS CLASSIFICATION	7
3	REQUIREMENTS RESOLUTION	10
3.1	User Support Activities and Services Requirements	10
3.2	Middleware Requirements.....	10
3.3	Operational Requirements.....	10
4	CONCLUSIONS	12
5	REFERENCES	14



1 INTRODUCTION

This document provides an analysis of the requirements gathered from user communities by the User Community Support Team (UCST) during the May-November period through the NA3 processes outlined in MS305 [R 1]. The requirements, or needs, as expressed by filling out the forms provided by NA3 span the whole spectrum of typical user needs, from training and documentation material, touching clearly identifiable middleware components and capabilities, and ending with the description of concrete bugs and non-functional demands.

The primary targets of this technical note are the members of the TCB, notably the Technology Providers, as it provides a prioritisation and categorisation of requirements (with respect to the EGI Capabilities) from the EGI User Community that should be implemented in any upcoming software release from the EGI Technology Providers.

2 REQUIREMENTS CLASSIFICATION

The requirements collected through the NA3 activity and aggregated into the requirements descriptions in the appendices of MS305 (version 6 was used) forms the basis for the classification provided in this section.

MS305 describes various channels of information gathering, such as surveys, direct unsolicited (but welcome) input and documents publicised and endorsed by other projects and communities (e.g. the ESFRI projects). Passive information gathering, such as document analysis, often provides information covering a constricted set of topics that regularly do not entirely match the structures and processes set up within EGI-InSPIRE. Active information gathering activities such as surveys, online forms of attendance at project and NGI events often provide intrinsic structure to the requirements covered or described in such activities.

Also, the level of detail provided in the collected requirements varies greatly. Data collected covers descriptions of bugs in a particular software component, the general failure of the implementation of a feature. It also covers clear demands to introduce named components (such as STORM) into the production infrastructure. The data also clearly describes non-functional requirements for a spectrum of components encountered in the middleware. To aid in the analysis of these requirements they have been classified into the different capabilities described in the UMD Roadmap. Requirements recorded in MS305 that do not relate to technology issues are not included in this document.

Requirements related to middleware that could not be associated with a specific EGI Capability, or were explicitly described for all known/available services to the VO are classified as cross cutting requirements, applicable to all EGI Capabilities.

The last two columns of Table 1 denote whether the described requirement is considered a functional (“f”), or a non-functional (“nf”) requirement in the software engineering sense.

Table 1: Categorisation of requirements

UMD Area	Capability	Requirement	Requestor	Ref.	f	nf
Functional	Compute	Support for pilot jobs	HUC/LifeSciences	p65	X	
		JDL misses GLUE attribute for available disk space on worker node	lofar, lsgrid	p79, p81	X	
	Compute Job Scheduling	Resubmit jobs that are stuck in the queue	cdf	p71	X	
		CREAM-CE stability	hone	p76		X
		Instability, scalability and performance issues of WMS	calice, ilc, fusion, lofar, lsgrid	p70, p77, p76, p79, p81		X
	File Access	LFC Admin: Management functions and procedures	HUC/LifeSciences	p65	X	
		DMS Workload management	HUC/LifeSciences	p65	X	
		Handling many small files instead of few large ones	HUC/LifeSciences	p65	X	
		Data lifetime management	vlemed	p84	X	

		SRM performance better than 2 requests per second	lofar, lsgrid	p79, p81	X
		LFC reliability, reduce orphan files and stale handles	lofar, lsgrid	p79, p81	X
	Parallel Job	Improved MPI-START support	complex-systems, voce	p85, p87	X
		Hybrid parallel jobs (MPI/OpenMP)	complex-systems, voce	p85, p87	X
		Passwordless ssh connection between worker nodes of parallel job	voce	p87	X
	Workflow	Parametric jobs, DAG-style job graphs	complex-systems	p85	X
Security	AuthN	Federated identity/sign-on, SSO	HUC/LifeSciences, enmr.eu	p65, p73	X
		X.509 certs infrastructural, not end user visible	HUC/LifeSciences, enmr.eu	p65, p73	X
	AuthZ	Fine-grained access control	vlemed	p84	X
		Access control propagation and enforcement (e.g. from Metadata catalogue to file access service)	vlemed	p84	X
Operational	VM Image Management	User-defined VM image execution	iscpif	p86	X
	VM Management	Management of User VMs (Not explicitly stated. Derived from VM Image Mgmt requirement)		p86	X
Other	Cross-cutting requirements	Coherent and homogeneous API for MW services	HUC/LifeSciences, vlemed	p65, p84	X
		Preferences for bug fixes over new features	desktopgrid.edges	p72	X
		More flexible configuration possibilities in YAIM	desktopgrid.edges	p72	X
		Improve error messages and error statuses on all services	ilc, lofar, lsgrid	p77, p79, p81	X
		Better integration of infrastructures based on UNICORE, ARC and on gLite so that the same client tools can be used with any resource	lldg, NorduGrid	p78, p82	X
		Coherent command line parameter design	pheno	p83	X
		Stability of command line parameter names and design	pheno	p83	X
		Services must not be single points of failures	HUC/LifeSciences	p65	X
		Service stability	fusion, hone	p74,	X

In addition, the requirements also contained a number of specific bugs that will be entered into the EGI Helpdesk for resolution if they are not already recorded. These are described below for information:

UMD Area	Capability	Requirement	Requestor	Ref.	f	nf
Functional		Implementation of GLUE attribute "LogicalCPUs is faulty: Claiming all logical CPUs of a worker node causes the job to never execute	lofar	p79	X	
Other	Cross-cutting requirements	YAIM update and config overwrites custom settings (p72)	desktopgrid.edges	p72	X	

Table 2: Bugs reported against Middleware

3 REQUIREMENTS RESOLUTION

The requirements captured in MS305 can be placed in three categories:

1. User Support activities and services requirements
2. Middleware requirements
3. Operational requirements

Each category requires a different approach to resolving the stated requirements.

3.1 User Support Activities and Services Requirements

This category includes those requirements that the UCST can resolve within the EGI-InSPIRE NA3 Work Package. It includes issues relating to enhancements and additions to the User Support services and activities provided within NA3. Any other, non-technology related requirement will be forwarded by the UCST to the appropriate channels or declined with the originator informed of the actions if appropriate.

Several requirements in this area have been captured in MS305 but are outside the scope of the target audience of this document (the TCB) and are not included here.

3.2 Middleware Requirements

The middleware related requirements cannot be resolved without involving technology providers through the TCB and the integration of these requirements into their release plans. These requirements focus primarily on the deployed middleware now originating from the EMI project.

The UCB identified and endorsed the following five requirements (in no particular order) as high priority for EMI to integrate into their development roadmap:

- Increased stability and scalability for gLite WMS
- Better (more verbose and informative) error messages
- Fixing the known bugs before adding new features
- Coherency of command line commands, parameters and APIs
- Better feedback about jobs, automated resubmission of jobs that are stuck on sites

The categorisation of the data in Table 1, reinforces the requirements endorsed by the UCB but also brings forward cross-cutting issues around Parallel Jobs and Authentication. These additional issues are:

- Parallel job support (MPI, OpenMP, MPI-START)
- Federated identity with Single Sign On (SSO)
- Backend use of X.509 (i.e. not user facing) for identity management

3.3 Operational Requirements

Some requirements effectively relate to the operational deployment of the technology coming from the provider, i.e. the underlying software needs to have both the required functionality and be deployed and used in conjunction with other deployed systems. These issues go beyond just functional technology and involve the agreement of policies between different providers. The UCB endorsed the following broader operational issues that needed urgent attention:

- Simplifying the authentication mechanisms and processes that are exposed to end users



- Improved reliability of middleware services with guaranteed quality of services, especially when submitting large number of jobs, and/or transferring large amount of data, large number of files. These two items are also prominent in the summary presented in Table 1.

4 CONCLUSIONS

The most frequent mentioned EGI Capabilities in the recorded requirements are:

- Compute Job Scheduling
- Parallel Job
- Authentication (some are cross cutting requirements)

However, when analysing which EGI Capabilities drew the most attention or coverage by from within the recorded requirements, the following EGI Capabilities are indirectly mentioned the most:

- File Access Capability
- Compute
- Compute Job Scheduling
- Parallel Job
- Authentication
- Authorisation

Reassuringly, the most frequently mentioned EGI Capabilities include those for which the most popular requirements were identified. Interestingly enough, the most popular EGI Capability, *File Access*, shows at the same time the largest variance with six requirements, each mentioned only once.

Combining the most popular requirements and most popular EGI Capabilities leads to the following capabilities as being most in demand by the user capability.

- Parallel Job
- Compute Job Scheduling
- Authentication
- File Access

Combing this analysis with the endorsements of the UCB to improve the characteristics of the delivered software rather than specific functionality, the focus of the technology providers must be to deliver:

- **Better (more verbose and informative) error messages**
- **Fixing the known bugs before adding new features**
- **Coherency of command line commands, parameters and APIs**

Across the middleware used within the Authentication, Parallel Job, Compute Job Scheduling and File Access capabilities. For job related capabilities there must be better feedback about the job status and the automated resubmission of jobs that are stuck on sites, the specifically the stability and scalability of the gLite WMS must be increased.

The EGI expects that EMI – together with the UCST and the user communities – defines measurable goals for these priorities and defines a development plan for EMI 1.0 (the next release) due in April 2011. The focus for this release is believed to be service enhancements (security, consolidation of services and libraries, usability) for ARC, gLite, UNICORE and dCache.

In addition a roadmap needs to be defined by EGI, in conjunction with the TCB and OMB, to deliver to the UCB a plan to migrate to:



- Simpler authentication mechanisms and processes so that the technical complexity of X.509 certificates are hidden from the users and that they can be obtained through federated identity mechanisms.
- Improving the reliability of middleware services with guaranteed quality of services, especially when submitting large number of jobs, and/or transferring large amount of data, large number of files.

It is expected that the impact of this analysis and progress on the issues it identifies will be reported at regular intervals to the UCB and at the next F2F meeting that will be held in conjunction with the EGI User Forum, 11-15th April 2011. In the meantime the UCST will work together with the user communities and with the middleware providers on expanding these requirements as required.



5 REFERENCES

R 1	MS305: User Feedback and Recommendations, (version 6): https://documents.egi.eu/document/211
R 2	D5.1: UMD Roadmap, final version: https://documents.egi.eu/document/100