

EGI-InSPIRE

VO PORTAL POLICY

Document identifier	EGI-SPG-VOPortal-V1_0
Document Link	https://documents.egi.eu/document/80
Last Modified	13/07/2010
Version	1.0
Policy Group Acronym	SPG
Policy Group Name	Security Policy Group
Contact Person	David Kelsey/STFC, UK
Document Status	Submitted
Approved by	Body who approved the doc
Approved Date	DD/MM/YYYY

Policy Statement

This Policy applies to all Portals operated by Virtual Organisations that participate in the Grid.



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, 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

	Body	Date
Approved by	EGI.eu Executive Board	DD/MM/YYYY

III. DOCUMENT LOG

Issue	Date	Comment	Author/Organization
1.0	13/07/2010	Imported from JSPG policy document with the same title. The only change to wording made was to update a link (URL). See https://edms.cern.ch/document/972973 (V3.2a, dated 6 Aug 2009) for the old JSPG document.	David Kelsey/STFC
2.0			
3.0			
4.0			



IV. 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 the 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 the ESFRI projects. 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 – structured international user communities – that are grouped into specific research domains. VRCs are formally represented within EGI at both a technical and strategic level.



TABLE OF CONTENTS

1 VO Portal Policy - Preamble	5
1.1 Definitions of Terms	5
1.2 Portal Classes	6
2 General Conditions	7
3 Special Conditions	7
3.1 “Closed Self-Contained Simple One-Click” Portals	7
3.2 “Parameter” Portals	8
3.3 “Data Processing” Portals	8
3.4 “Job Management” Portals	8

1 VO PORTAL POLICY - PREAMBLE

This Policy applies to all Portals operated by Virtual Organisations that participate in the Grid.

1.1 Definitions of Terms

Portal

a web site or web service that provides functionality to Web Users via web-specific applications.

Web User

a human individual that accesses Grid resources through a Portal. This individual may or may not be (also) enrolled in a Virtual Organisation

Grid User

a human individual registered in a Virtual Organisation

Anonymous Web User

a Web User who does not provide unique credentials to the Portal when invoking functionality

Pseudonymous Web User

a verifiably-human Web User that provides authenticated non-identifying information to the Portal when invoking functionality. The aim of verifying that the user is a human is to prevent "automated" use of the portal to stop overload of the portal or use by another service. There are several ways that this could be achieved, e.g. a captcha, a one-time email address on a non-authenticated email (gmail, hotmail, etc) or knowledge that the portal can only be used by people sitting at a public login station (e.g. library walk-up system).

Identified Web User

a Web User that provides authenticated personal identification to the Portal when invoking functionality, but whose credentials and way of authentication are not necessarily compatible or equivalent with Grid authentication.

Strongly Identified Web User

a Grid User that provides authenticated identification to the Portal when invoking functionality, that allows the portal to authenticate to the Grid Resources with valid Grid credentials specific to the Grid User.

Resource Provider

a Grid participant that provides compute, data storage or database resources

Robot

a software agent that performs automatic functions on behalf of a natural person.

Secured Robot

a Robot whose private key was generated on and is held in plain-text exclusively on a Secure Hardware Token as defined by the IGTF in the ISCP urn:oid:1.2.840.113612.5.2.3.1.1 (see <http://www.eugridpma.org/guidelines/lscp/lscp-private-key-hardwaretoken-1.1.pdf>).

Robot Certificate

a certificate issued to a non-human entity actively acting as an automated client towards other entities as defined by the IGTF in the ISCP 1.2.840.113612.5.2.3.3.1 (see <http://www.eugridpma.org/guidelines/lscp/lscp-certtype-robot-0.1.pdf>). The common name in the robot certificates shall identify at least the natural person or group of persons responsible for the Robot. A Grid may mandate the use of a Secure Robot.

1.2 Portal Classes

Portal Classes			
Portal Class	Executable	Parameters	Input
Simple one-click	provided by portal	provided by portal	provided by portal
Parameter	provided by portal	chosen from enumerable and limited set	chosen from repository vetted by the portal
Data processing	provided by portal	chosen from enumerable and limited set	provided by user
Job management	provided by user	provided by user	provided by user

This Policy applies to Portals that belong exclusively to one of the following classifications:

Simple one-click portals

The Web User invokes functionality on the Portal where jobs submitted to the Grid use executable code that is provided by the Portal to the Grid as part of the job submission process. All parameters and input data are defined exclusively by the Portal and cannot be influenced by the user.

Parameter portals

The Web User invokes functionality on the Portal where jobs submitted to the Grid use executable code that is provided by the Portal to the Grid as part of the job submission process. The Web User may only provide run-time parameter settings from an enumerable and limitative set, and may select data files from a enumerable repository of data files that are pre-vetted for use by the Portal.

Data processing portals

The Web User invokes functionality on the Portal where jobs submitted to the Grid use executable code that is provided by the Portal to the Grid as part of the job submission process. The Web User may provide run-time parameter settings from an enumerable and limitative set, and may provide non-validated input data to the executable code, but where the user cannot influence the instructions executed.

Job Management portals

The Web User invokes functionality on the Portal where jobs submitted to the Grid use executable code that is provided by the Web User. Whether this code is passed through unmodified by the Portal and is submitted to the Grid as-is, or whether this code is inspected and analysed on the Portal does not change the classification of this Portal.

Frameworks that submit jobs to the grid for multiple users based on a late-binding mechanism are subject to the ‘security policy for operation of multi-user pilot jobs’ and are not covered in this Policy.

Portals that cannot be classified within the framework laid down by this policy should be reviewed independently, in accordance with the spirit and intention of this Policy, after such an evaluation this policy should be amended with this portal class, or they should be classified as a Job Management portal.

Portals whose functionality fits multiple classes must ensure they comply with the conditions that are applicable to the currently selected functionality.

2 GENERAL CONDITIONS

All Portals, operated by or on behalf of a Virtual Organisation, must comply with the Virtual Organisation Operations Policy.

In addition to all other Policies, the following conditions apply to all Portals:

- The Portal, the VO to which the Portal is associated, the Portal manager are all individually and collectively responsible and accountable for all interactions with the Grid, unless a credential of a Strongly Identified Web User is used to interact with the Grid.
- The Portal must be capable of limiting the job submission rate.
- The Portal must keep audit logs for all interactions with the Grid as defined in the Traceability and Logging Policy (<https://documents.egi.eu/document/81>).
- The Portal manager and operators must assist in security incident investigations by the Grid and by any Resource Provider who can justifiably claim to provide or to have provided resources to the Portal.
- Where relevant, private keys associated with (proxy) certificates must not be transferred across a network, not even in encrypted form. Other re-useable password or private key data should not be transferred across a network, but if transferred must be encrypted when sent across a network.
- The Portal must not persistently store passwords or private keys for its end-users that can be used to authenticate to the Grid past 1 million seconds. This is aligned with the definition of "short-lived" authentication credentials used on the Grids.
- When a Portal Credential is used to store data on the Grid as a result of an action by a User, it may only be stored in locations that have been specifically agreed between the Portal and designated Resource Providers and only for as long as the User is associated with the portal. Transient data may be stored in designated scratch areas on the computational resources provided to the running jobs. When a Grid User Credential is used, data may be stored in all Grid locations where the Grid User has permission to store such data.

3 SPECIAL CONDITIONS

Depending on the Portal class, the following conditions will specifically apply.

3.1 *"Closed Self-Contained Simple One-Click" Portals*

By registering a Closed self-contained simple one-click Portal in a Virtual Organisation, you agree to the conditions laid down in this section and documents references therein.

1. The Portal may offer services to all Web Users.
2. The Portal must use a Robot Certificate to interact with the Grid.
3. Maximum submission rate must be specifically agreed between Portal and Grid
4. The Portal must keep enough information to associate any interactions with the Grid with a particular Internet address and (tcp) port used by the requester.

3.2 “Parameter” Portals

By registering a Parameter Portal in a Virtual Organisation, you agree to the conditions laid down in this section and documents references therein.

1. The Portal may offer services to Pseudonymous, Identified or Strongly Identified Web Users.
2. The Portal may either use a Robot Certificate, or use the Grid credential for Strongly Identified Web Users.
3. The job submission rate may be limited differently for Pseudonymous, Identified and Strongly Identified Web Users, and the maximum submission rate by the Portal induced by Pseudonymous and Identified Web Users must be specifically agreed between you and the Grid.
4. The Portal must keep enough information to associate any interactions with the Grid with a particular user. If the user was Identified or Strongly Identified, relevant authentication information must be recorded and archived.

3.3 “Data Processing” Portals

By registering a Data Processing Portal in a Virtual Organisation, or by connecting a Data Processing Portal to the Grid infrastructure, you agree to the conditions laid down in this section and documents references therein.

1. The Portal may offer services to Identified or Strongly Identified Web Users.
2. The Portal may use a Robot Certificate, or use the Grid credential for Strongly Identified Web Users.
3. The Portal must keep enough information to associate any interactions with the Grid with a particular Web User. Relevant authentication information must be recorded and archived.
4. The system used to authenticate Identified Users must be adequately secured. In particular additional requirements apply:
 1. Web Users must be notified of all registrations, modifications and of removal of their data in the authentication database.
 2. The authentication database must contain enough information to contact the Web User for as long as the user is registered.
 3. Entering authenticating information in the database, including resets of such information, must be appropriately authenticated.

3.4 “Job Management” Portals

By connecting a Job Management Portal to the Grid, you agree to the conditions laid down in this section and documents references therein.

1. The Portal may offer services only to Strongly Identified Web Users.
2. The Portal must use Grid User credentials specific to the Web User and use these for all interactions with the Grid.
3. The Portal operations must comply with the Site Operations Policy.