



EGI-InSPIRE

VO REGISTRATION PROCESS IN EGI

Document identifier:	EGI-VO_registration_process_v4
Date:	23/01/2012
Activity:	NA3
Lead Partner:	EGI.eu
Document Status:	FINAL
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/278

Abstract

The document describes the process of enabling a Virtual Organisation (VO) on the European Grid Infrastructure (EGI) and the parties who are involved in process execution. Users of EGI are organised into Virtual Organisations (VO). A VO is a group of people (typically application scientists and application developers) who share similar interests and have similar goals and who need to work collaboratively and/or need to share resources (e.g. data, software, expertise, CPU, storage space) through a grid infrastructure regardless of their geographical location. The focus of this document is on the tasks that VO representatives and the EGI staff have to accomplish in order to register and validate a new VO on EGI. The goal is to capture the VO registration workflow so it can be executed and further improved by user communities and by the EGI staff. At the moment of writing this workflow is closely coupled to the implementation of the Operations Portal (<https://operations-portal.egi.eu>).



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. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
0.1	14/07/2010	First draft	Gergely Sipos, EGI.eu
0.2	16/07/2010	Input for second draft	Rolf Rumler, IN2P3
0.3	20/07/2010	Third draft	Gergely Sipos, EGI.eu
0.4	15/09/2010	Fourth draft with inputs from Rolf Rumler and Cyril Lorphelin	Gergely Sipos, EGI.eu
0.5	30/09/2010	Fifth draft	Goncalo Borges, LIP
0.6	18/11/2010	Applied EGI.eu document template, removed responses from Cyril as they are recorded in the requests document	Karolis Eigelis, EGI.eu
v1	10/12/2010	First public draft	Gergely Sipos, EGI.eu
v2	10/12/2010	DocDB link inserted	Gergely Sipos, EGI.eu
V3	10/12/2010	Formatting	Gergely Sipos, EGI.eu
V4	26/05/2011	Updating based on new operations portal release 2.6.1	Karolis Eigelis, EGI.eu

III. APPLICATION AREA

The document is targeted to all stakeholders and collaborators of EGI who want to learn about the VO registration process. The document captures the registration process as it is controlled by the Operations Portal, available at <https://operations-portal.egi.eu>.

IV. DOCUMENT AMENDMENT PROCEDURE

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

V. TERMINOLOGY



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



VI. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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



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

VII. EXECUTIVE SUMMARY

In the context of this document a Virtual Organisation is defined as a group of researchers and application developers with similar interests and requirements, who are able to work collaboratively regardless of geographical location, using a grid infrastructure. The European Grid Infrastructure and the EGI-InSPIRE project uses the Operations Portal [R6] to keep track of VOs of the infrastructure, to provide information about VOs to EGI members and the generic public and to register and validate new VOs.

The focus of this document is on the tasks that need to be performed by a VO and by EGI staff during the VO registration process. The purpose of the document is to capture the VO registration workflow so it can be learned by VO representatives, by EGI staff as well as it can be improved in order to meet new requirements.

For other aspects of VO management (e.g. operation support, resource/service allocation, decommissioning) please consult with the VO services Wiki page [R7] or contact the VO services team via EGI Helpdesk [R8].

The process that is captured in this document is based on a similar process that EGI-InSPIRE inherited from the Enabling Grid for E-science (EGEE) project and which is specified in [R4]. The EGI VO registration workflow includes changes compared to the EGEE version, mainly, because some of the stakeholders of the EGEE process do not exist in EGI.



TABLE OF CONTENTS

1 THE VO REGISTRATION PROCESS.....	7
1.1 Accepting a VO	9
1.1.1 Valid VO Id cards.....	9
1.1.2 Scope of the VO	12
1.1.3 VO lifecycle – VO states	12
1.2 Rejecting a VO registration request	13
2 SUPPORT FOR VO OPERATION AND MONITORING	14
3 FEEDBACK AND SUGGESTIONS.....	15
4 REFERENCES	16

1 THE VO REGISTRATION PROCESS

Any person, who holds a grid certificate recognised by EGI [R9]. can register a new EGI VO via the Operations Portal. The person who initiates the registration is called the *VO manager*. After the VO is setup and operational, the VO manager is the person who is primarily responsible for the operation of the VO and for providing sufficient information about VO activities for EGI and for VO members (to both people and sites).

The registration process inside EGI is handled by and coordinated by the *VO supervisor*. This person is delegated from the User Community Support Team of EGI.eu (UCST).

The following table describes the VO registration process, listing each of the steps that need to be performed, the people who are responsible for the action, and the physical action that need to be executed to complete the step.

#	Step	Performed by		Tool to use	Physical action	Prerequisites (if any)
		External to EGI	Internal to EGI			
1	Submit VO registration request	VO manager		Operations portal	Fill out Web form (the VO Id card) [R1]	Grid certificate must be in Web browser
2a	Inform EGI about new VO registration request		Operations portal		Send notification email to <ul style="list-style-type: none"> •VO supervisor¹ •NOC/ROC managers list 	
2b	Inform EGI that a new VO requires VOMS server		Operations portal		Open GGUS ticket requesting a VOMS server to the new VO, and asking to be assigned to the VO Services support unit.	VO manager asked EGI for a VOMS server in Step 1.
2c	Inform EGI that new VO requires new Support Unit in GGUS		Operations portal		Open a ticket against the GGUS support unit in GGUS.	VO manager asked assistance in Step 1 to setup a new GGUS Support Unit.
3a	Check the correctness of VO registration request		VO supervisor	Operations portal	Check the content of the VO Id card [R2]	<ul style="list-style-type: none"> •Grid certificate must be in Web browser •Must have "OAG manager" role in operations portal

¹ "VO supervisor" is the member of the EGI User Community Support Team who is responsible for the approval of VO registration requests. The VO supervisor can be contacted through the ucst@egi.eu email address.

3b	Ask for update of VO Id card		VO supervisor	GGUS ticket	Send out update request to VO manager	Data is missing or incorrect in VO Id card ²
4	Recommend the VO for approval		VO supervisor	GGUS ticket	If the scope of the VO is GLOBAL, then send approval recommendation to CCO and COO. They have 5 work days to raise any objection ³ .	VO Id card contains correct and complete data
5	Approve registration		VO supervisor	Operations portal	On the VO Id card [R3]: •Set VO status from NEW to VALIDATED •Set the scope of the VO Save the Id card	The scope of the VO is GLOBAL and there is no objection against approval from CCO and COO.
6	<i>Optional step:</i> Setup VOMS server and register in GOC DB.	VO manager		VOMS specific	•If this step is not included, then the VO must ask EGI for a VOMS server in step 1. •If this step is included, then it can happen here, or even before step 1.	The VO cannot use any of the existing VOMS servers.
7	Defining VOMS server on VO Id card	VO manager		Operations portal	Can be part of step 1.	
8a	Inform EGI that Id card contains VOMS server		Operations portal	Email	Send notification email to •VO supervisor •NOC/ROC managers list •Email list of "VO services" group of EGI	
8b	Inform EGI that New VO SU was created in GGUS		Operations portal	Email	Send notification email to •VO supervisor •NOC/ROC managers list	
9	Approve new VO Id card		VO supervisor	Operations portal	Set VO status from PENDING or NEW to ACTIVE, save VO Id card in [R3].	
10	Inform EGI about new		Operations portal	Email	Send notification email to •VO supervisor	

² See the specification of correct and complete VO Id cards in Section 1.1.

³ CCO: Chief Customer Officer; COO: Chief Operations Officer

	ACTIVE VO				•NOC/ROC managers list	
11	NOC/ROC managers invite sites to join the VO		NOC/ROC managers	Email	Forward email to site administrators	
12	Sites join the VO		EGI Operations	Internal to EGI Operations		

1.1 Accepting a VO

Before a VO can be accepted the VO supervisor must:

1. Verify that there is no existing VO with significantly overlapping goals. This can be done through the VO list of the Operations portal [R5]. VOs with similar goals (e.g. image analysis) should be advised to join.
2. Check that the VO Id card contains *correct and complete* data.
3. Choose the scope of the VO.

1.1.1 Valid VO Id cards

The following compulsory and optional fields must be filled out by the VO manager as part of the registration process (Step 1 in the table above):

•Section **General information**

○Mandatory field **Name**

The Operations portal enforces a DNS style name. It still has to be verified whether the VO manager whose name and mail address is available in the Contact list update section is authorised to use it. The VO registration procedure requests this but currently no enforcement is done.

It is checked, though, whether the VO name is already in use, and if so, portal pops up the notification asking to choose another name. The obtained information is given back to the VO manager if it is not obvious that the owner of the domain and the VO manager are the same person. Note that it is not considered sufficient that the VO manager's mail address is in the same domain as the VO name's one, nor that the VOMS server or VO home page address are of that domain, if this information is available.

Doubts on domain ownership are not stopping VO registration, as the responsibility of acquiring the domain name is with the VO manager anyway.

○Mandatory field **Description**

In principle any text is valid. However, it should describe a scientific or technical activity, or should be related to education. The text is also used to delimit proper resource usage on the grid, so it should be significant for this purpose, i. e. saying

“VO giving access to the grid” is a poor description whereas “VO giving access to the grid for training purposes” is completely satisfying.

In practice up to now every VO request came with a readable text but some VOs got stuck in the very first stage of the registration (state NEW) because of a too minimalistic view of what is a description.

○Mandatory field **Discipline**

It is simply verified whether there is a contradiction between the field Description just discussed and this one.

○Mandatory field **Supported Middleware**

There are four options to choose which middleware the VO support, portal automatically checks that at least one option was chosen by Vo Manager.

○Mandatory field **Acceptable Use Policy (AUP)**

The acceptable use policy which is meant here is the VO AUP.

On the “New VO registration web page” [R 6] the registering VO manager has the choice between a text automatically generated from the Description but where at least some words have to be updated, or a file in text or pdf format uploaded by the manager containing a VO written AUP.

In the former case it has just to be checked whether the update has been done; the words to be replaced are “owner body”, included in brackets - “[]” - , and the replacing text must specify the authority enforcing the VO AUP. This is however omitted in one out of two cases but then normally corrected rapidly by the VO manager. If not the VO gets stuck in the NEW state; there are still some of them.

If the AUP is uploaded, the complete text has to be verified if it corresponds to a VO AUP. In case of a doubt, in addition to contacting the VO manager a member of the JSPG is asked for advice.

○Mandatory field **VO homepage**

This field must be verified whether the home page contains information about the ongoing/planned grid activity and that this information corresponds to the VO’s Description. Sometimes the scope of the VO can also be determined with this or with the VO manager’s affiliation. (For example about the scientific goals of the community and how the EGI VO helps the community to achieve these goals.)

○Mandatory field **Enrolment URL**

This field must be verified whether it is functional or simply an optional service to the VO. Additionally, the information available on the enrolment web page *might* give some indications on the purpose and scope of the VO as well as on the attitude concerning security (availability of a Grid AUP, reminder of correct resource usage etc.).

- Section **VOMS Information**

- Mandatory field “**VOMS Configuration**”

- There are two options the VO Manager must choose: one is a VOMS server which is pulled from GOCDB and another is a request for support in setting up the VOMS server.

- Section **VO SU at GGUS**

- Optional field “**check box**”

- There are two options the VO Manager can choose: one is a default – No. If VO Manager will check a box, the new ticket will be created for GGUS support unit and VO Supervisor will keep track of the process.

- Section **Generic Contacts**

- There is only one not mandatory contact in the list of this section shown on the VO ID card, **Operations contact**. Other fields (VO Managers, Security, User Support, VO Users) are mandatory and currently, new registration requests must contain a valid address in these fields. Validity should be checked by sending an e-mail to it, requesting confirmation of receipt.

- Section **Change status & scope**

- Pull down list **Scope**

- As already indicated in the discussion of the previous fields, any hints are used to determine the value to be selected for Scope. In case of a doubt - which is the normal case here - a suggestion is made to the VO manager. The field is then updated only after a feedback from that person.

- Assigning a correct value is important for limiting the noise especially on the NOC/ROC managers list in case of Regional VOs and also to determine responsibilities for support in case of additional resource requests made by the new VO.

- If the VO is a Regional one, this field should be updated **before** the Status field.

- Updating this field triggers notifications to the VO Services group list and to the NOC/ROC managers list in all cases.

- Pull down list **Status**

- If all previously mentioned fields contain valid values, either since the beginning or after some communication with the VO manager, the status can be changed from **NEW** to **PRODUCTION**.

- The VO will be then active and in production state.

- Notifications are sent to to the VO Service group list in all cases and to the NOC/ROC managers list in all cases except for Regional VOs where only the corresponding NOC/ROC is informed.



1.1.2 Scope of the VO

As part of the VO approval step (Step 5 in the table above) the scope of the VO must be defined by the VO supervisor based on information provided by the VO manager either in the VO Id card, or through additional channels (e.g. in email). The scope must be one of the following:

- **GLOBAL:** the VO is supported by sites from multiple countries and all of these countries are represented by its National Grid Infrastructure (NGI); comprises an international user community and/or has international resources coming from sites of different countries represented by their National Grid Infrastructures (NGIs).
- **NATIONAL:** at least the supporting sites of the VO belong **to only one NGI**; i.e. sites and users are located within the same country. Users *might* come from elsewhere but they are working inside the scope of the same NGI where the sites are. The associated NGI is part of the scope, like for example “NGI - Italy” or “NGI - France”.

In case of invalid, unclear or ambiguous entries in any of the controlled fields of the VO Id card, or in case of doubts about the goals of the VO, the requestor must be contacted and invited to clarify the situation or to correct the entries.

1.1.3 VO lifecycle – VO states

A VO is in one of the following states:

- **NEW:** this is the initial state when the VO creation is requested. It is automatically assigned.
- **PRODUCTION:** the “normal” state of a VO. It is manually given to a VO by the VO supervisor when the VO manager has entered a valid VOMS server on the Id card of the VO.
- **SUSPENDED:** this state is entered when the VO no longer has a VOMS server, either because (one of) the VO manager(s) deleted the corresponding entry on the VO ID card, or some other person with VO Supervisor role in the operations portal did that. This may be temporary or the preparation of VO deregistration. Manual intervention is needed to put a VO into this state.
- **DELETED:** the final state of all VOs ever registered. VOs where the registration was rejected do not get here. No trace is kept of erroneous registration requests neither. The associated VO Id card should basically be empty, the only information which really can be considered valid is the VO name (and the state itself). Keeping the VO name is meant to avoid giving the same name to different VOs and so to avoid confusion for sites which forgot or considered unnecessary to suppress a DEAD VO from their information system. It also keeps historical accounting data consistent.

Note that manual state changes can only be made by people registered in VO Supervisor role on the Operations portal or by the Operations portal team itself. This document covers the VO lifecycle from “non existing” through NEW, then VALIDATED to ACTIVE.



1.2 Rejecting a VO registration request

Before the Operations portal enforced the DNS naming rule automatically, one request was rejected because it did not respect this rule. Later on at various times VOs were registered with a valid DNS like VO name and then the VO manager asked to change it to a non DNS one, or the EGEE RAG was asked before registration whether a non DNS name would be possible. Such requests for VOs which were not deployed on the EGEE infrastructure before the DNS naming rule was enforced were rejected. No other rejections took place.

However, several VOs got stuck in the registration process because they did not correct some invalid mandatory information, especially in the VO AUP, the VO Description, or the VO Homepage despite of repeated reminders to do so.



2 SUPPORT FOR VO OPERATION AND MONITORING

- If the scope of a VO is national (i.e. both users and sites belong to a single country) then support for the VO must be provided by the respective National Grid Initiative/Infrastructure (NGI). The User Community Support Team can connect the VO manager to the respective NGI support team (ucst@egi.eu).
- The “VO Services” team of EGI provides assistance for the setup and operation of VOs. The team is able to answer questions related to VO operation, monitoring and accounting. The team also provides documentation and catch-all services for VOs. These catch-all services cover the basic monitoring and accounting needs of most communities. The VO Services team maintains up to date information at its Wiki page [R7] and can be contacted via the EGI Helpdesk helpdesk: www.ggus.org (Name of the support unit is also “VO Services”).
- The VO supervisor can obtain an overview of VOs at the following page: <http://operations-portal.egi.eu/vo> This table indicates which fields of the VO Id cards are filled / missing by each of the VOs. The missing bits are highlighted in red. Green indicates properly filled segments.



3 FEEDBACK AND SUGGESTIONS

Feedback and suggestions relating the VO registration and support processes must be sent to the User Community Support Team of EGI.eu (ucst@egi.eu) or to the “VO Services” support unit in the EGI Helpdesk.

4 REFERENCES

R 1	VO registration page in the Operations portal: https://operations.gridops.org/index.php?section=vo&page=newvoregistration
R 2	VO request handling page in the Operations portal: https://operations.gridops.org/index.php?section=vo&page=idcardupdate
R 3	VO request approval page in the Operations portal: Can be accessed from [R2] by clicking on the name of the pending VO request.
R 4	EGEE workflow of VO registration (practical guide): https://edms.cern.ch/file/1070320/1/EGEE-III-WorkflowVOregistration-v0_1.doc
R 5	List of VOs in the Operations portal: http://operations.gridops.org/index.php?section=home&page=volist
R 6	Operations Portal: http://operations.gridops.org/
R 7	VO services Wiki page: https://wiki.egi.eu/wiki/VO_Services
R 8	VO Services Support Unit in the EGI Helpdesk: https://wiki.egi.eu/wiki/GGUS:VO_Services_FAQ
R 9	Grid certification authorities recognised by EGI: http://www.igtf.net/