



EGI-InSPIRE

RESOURCE ALLOCATION TASK FORCE

List of possible 3 QoS levels

Alessandro Costantini, Franck Michel, Marco Verlatto

<https://rt.egi.eu/rt/Ticket/Display.html?id=5061>

Date:	28/03/2013
Document Status:	Final
Dissemination Level:	Internal
Document Link:	https://documents.egi.eu/document/xxxx

Abstract

This document provides a list of possible 3 QoS levels which can be proposed to Resource Providers (RP) on the bases of the action on Alessandro Costantini, Franck Michel, Marco Verlatto as members of the Resource Allocation Task Force of EGI.



COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. 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, 2010-2014. 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.

I. DOCUMENT LOG

Version	Date	Comment	Author/Partner
1.0	14/03/2013	Draft	A. Costantini
1.0	21/03/2013	Draft	F. Michel
1.0	21/03/2013	Draft	A. Costantini
1.0	26/03/2013	Draft	A. Costantini
1.0	27/03/2013	Draft	F. Michel
1.0	27/03/2013	Draft	M. Verlato
2.0	28/03/2013	Draft	A. Costantini
2.0	28/03/2013	Draft	F. Michel
3.0	28/03/2013	Final	A. Costantini

II. ACKNOWLEDGMENT

None

III. APPLICATION AREA

This document is a public report produced by the members of the “Resource Allocation“ EGI Task Force, on behalf of the EGI Council. Further information is available at https://wiki.egi.eu/wiki/Resource_Allocation_Task_Force

IV. TERMINOLOGY

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



TABLE OF CONTENTS

RESOURCE ALLOCATION TASK FORCE	1
1 INTRODUCTION	4
2 MANDATE	4
3 ACTION	4
3.1 Mandatory	5
3.2 QoS list	5
3.2.1 Opportunistic (basic support level provided by a Pool).....	5
3.2.2 Fair Share.....	5
3.2.3 Reserved.....	6
4 CONCLUSIONS.....	6
APPENDIX A – VO REQUEST SUBMISSION TEMPLATE.....	6
ANNEX.....	7
REFERENCES.....	7



1 INTRODUCTION

EGI needs to develop a model and the related processes for allocation of resources as a service to international collaborations. Offering resources as a service is relevant to all communities that do not contribute own resources to the infrastructure.

Policies, a service model and the related processes are needed to:

- provide resources in a streamlined fashion making its procurement easy for users, while taking into account the different Resource Centre-level and NGI-level resource allocation policies
- provide resources as a service through service level agreements in order to give some structure to resource offering for VOs relying on opportunistic usage of resources
- ensure mechanisms are in place to acknowledge usage of resources by VOs that rely on opportunistic usage of the infrastructure. This can happen through acknowledgement of infrastructure usage in published research work and it is important for Resource Centres and NGIs to provide evidence of their contribution to national and international science
- develop processes to express resource demand, and resource offer.

In Nov 2012 the EGI Council approved the start of experimental activities around the development of policies and processes for the application, scientific review and pooled resource allocation of NGI resources. [More information](#)

This task force aims at focusing on the processes and the operational aspects of resource application and allocation. The process for assessment is out of scope and will be dealt with by other groups.

2 MANDATE

- Collect information about best practices and processes that are already in place at a national level
- Propose processes around central management of demand and offer of pooled NGI resources, through the establishment of
 - SLAs to express user demand (involved parties: customers, EGI.eu)
 - OLAs to express resource offer (involved parties: EGI.eu, Resource Providers)
- Identify the needs of development in the existing operational tools for the automation of the identified processes
- Coordinate the implementation of a testbed with the involvement of VOs and Resource Providers

3 ACTION

Action on Alessandro Costantini, Franck Michel, Marco Verlatto:

Create a list of possible 3 QoS levels which can be proposed to Resource Providers (RP). Address concerns of RP and users.

The list is supported by a template (VO Request Submission Template, APPENDIX A) that has to be filled-in by the VO.

3.1 *Mandatory*

Inputs: Material stored on the Indico page

<http://indico.egi.eu/indico/conferenceDisplay.py?confId=1346>

In order to define the 3 QoS levels these points have been assumed

1. the amount of resources allocated for each VO Request for Resources by each Resource Provider Pool (RPP) or NGI Pool (NGIP) has to be decided between the Operator and the pools, and formalized in the resource offer. Examples: not more than 10% of the resources of the RP should be allocated to a single VO; not more than 50% of a VO Request should be covered by the resources of the RP.
2. the authentication method used by each Pool is the one already in use in the EGI Infrastructure: certificate credential X509 based method;
3. the accounting and the management of the pool accounts is left to the Resource Provider admin of the Pool;

None of the above points affects the QoS level definition.

3.2 *QoS list*

3.2.1 **Opportunistic (basic support level provided by a Pool)**

- Resources not guaranteed and subject to availability: CPU slots are granted when no other job is submitted by a VO with a higher priority.
 - o The Pool supports the VO for a fixed period of time. The duration of the agreement has to be declared in the VO Request submitted to the Operator and is subject to negotiation
 - o Before expiration, the agreement can be renewed and renegotiated directly with the RPP or NGIP
 - o 3 different queue levels are proposed
 - short: 30 minutes max - for NAGIOS probes or short test
 - medium : 12 hours max
 - long : 48 hours max
- The opportunistic model is commonly the default model for RPs that wishes to support a VO without any formal agreement. Formalizing this model can help VOs to be more aware of who supports them, and may foster resource acknowledgement.

3.2.2 **Fair Share**

- Resources available in fair share mode for a fixed time period: during the period, the VO is allocated a maximum number of CPU hours per day/week/month; the day/week/month period is hereafter called Fine Time Scale (FTS). When the VO reaches this threshold, new jobs are rejected. The VO has to await the next FTS for jobs to be accepted again.
 - o The Pool supports the VO for a fixed period of time. The duration of the agreement has to be declared in the VO Request submitted to the Operator and is subject to negotiation
 - o After expiration a new Request has to be submitted to the Operator

- Fine Time Scale (FTS) may be declared in the VO Request submitted to the Operator, or may be proposed by the Pool operator, and is subject to negotiation by the Pool
- 3 different queue levels are proposed
 - short: 30 minutes max - for NAGIOS probes or short test
 - medium : 12 hours max
 - long : 48 hours max

3.2.3 Reserved

- resources are exclusively reserved to the VO
 - available at any time during a fixed period of time
 - 3 different queue levels are proposed
 - short: 30 minutes max - for NAGIOS probes or short test
 - medium : 12 hours max
 - long : 48 hours max
 - If limit is reached (period or wall time) new jobs cannot be submitted
- Would probably be appropriate for well-organized large communities that are able to make 100% use of the dedicated cores during the agreed period.

4 CONCLUSIONS

This document provides a list of possible 3 QoS levels which can be proposed to Resource Providers (RP) on the bases of the action on Alessandro Costantini, Franck Michel, Marco Verlatto as members of the Resource Allocation Task Force in EGI

The list is focused on the features of the 3 services that has to be provided by Resource Provides:

- Opportunistic. Resources not guaranteed and subject to availability
- Fair Share. Resources available in fair share mode for a fixed time period
- Reserved. Resources are exclusively reserved to the VO.

The “Reserved” model has to be deeply investigated in order to clarify some aspects mainly related to the two different policies for resource allocation:

- Walltime allocation (PL-Grid) for a given period
- CPUs allocation (others) of a given period.

The list is supported by a template (VO Request Submission Template, APPENDIX A) that has to be filled-in by the VO.

APPENDIX A – VO REQUEST SUBMISSION TEMPLATE

The VO should express his VO Request for Resource according to a template like this:

- Current status info (estimated for new VOs):
 - # of VO registered users (or estimated user potential)
 - Geographical distribution of the users (% EU, % North America, % Asia-Pacific, % Africa, % Latin America)

- # of RPs currently supporting the VO, if any (e.g. 24 opportunistic, 0 fair share, 3 reserved)
- Average or estimated Normalized CPU time/Wall time per month in the last year
- Average or estimated # of completed jobs per month in the last year
- Average CPU time/Wall time per job/application type, e.g. X hours per application type A, Y hours per application type B, etc.
- Request of additional resources to the RA Pool
 - Time frame of the request (renewable after the end)
 - Opportunistic model: as many as possible RPPs and NGIPs willing/interested in supporting our community. This request will probably be the same for all VOs, so maybe can be considered the default, so no need to specify it in the VO Request for Resource. From the Pool perspective however it could make sense that the fulfillment of this request is somehow in proportion of the geographical users' distribution, e.g. if most of the users of a VO are located in EU, NA and Asia-Pacific the baseline opportunistic support will be asked at first to RPPs and NGIPs located in EU, NA and Asia-Pacific areas.
 - Fair share model: my VO is asking for XXX normalized CPU hours and Wall time hours per day/week/month/year.
 - Reserved model: a fraction YY% of my XXX normalized CPU/Wall time hours would like to be run on reserved CPU cores, to ensure immediate access to a few of resources for testing purposes or more urgent tasks.

ANNEX

None

REFERENCES

[1] https://wiki.egi.eu/wiki/Resource_Allocation_Task_Force_brainstorming_page