EGI-InSPIRE

Annual Report on the EGI Production Infrastructure

EU DELIVERABLE: D4.8

Document identifier:	EGI-D4.8-FINAL
Date:	**14/05/2013**
Activity:	**SA1**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/1664

Abstract

This document provides information on the status of the EGI Resource Infrastructure at the end of PY3. In particular, it describes the status and progress of Resource Centres, Resource infrastructure Providers and Operations Centres that are responsible of the daily operations of the infrastructure used by the supported research communities. The document provides information on the amount of installed capacity provided, the status of the current EGI user base, the trends in usage, the service levels provided and the status of VO Services and grid common infrastructure services. The status of the Staged Rollout infrastructure for software testing is also presented.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	M. Krakowian	EGI.eu	26/04/2013
Reviewed by:	J. Wolfrat G. Taffoni S. Newhouse	SurfSARA INGV/SA3 EGI.eu	07/05/2013 07/05/2013 01/05/2013
Moderated by:	K. Eigelis	EGI.eu/NA2	08/05/2013
Approved by	**AMB & PMB**		13/05/2013

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	05-04-2013	First draft	T. Ferrari/EGI.eu M. Krakowian/EGI.eu
2	26-04-2013	V1.0 for external review	T. Ferrari/EGI.eu M. Krakowian/EGI.eu
3	01-05-2013	V1.1 incorporating S. Newhouse's comments	T. Ferrari/EGI.eu
4	07-05-2013	V1.2 incorporating reviewers' comments	T. Ferrari/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: https://wiki.egi.eu/wiki/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

During PY3 SA1 was responsible of the continued operation and expansion of the production infrastructure. The total number of Resource Centres (RCs) in March 2013 amounts to 340 certified instances (this figure does not include 27 instances that were temporarily suspended for maintenance or underperformance). The number of temporarily suspended RCs increased due to two grid middleware upgrade campaigns that were undertaken starting in PQ10 for the decommissioning of gLite and EMI-1 software. Other factors contributed to this reduction in number, like the end of operations of NGI Ireland and Iniciativa de Grid de America Latina – Caribe. In PY3 the integration of the Ukrainian National Grid was successfully finalized. In PY3 two new countries started contributing resources to the Asia Pacific federated Operations Centre: Iran and Vietnam

The decrease in the number of production RCs was however compensated by a substantial increase in the offered capacity: compute resources increased by +33.6% in PY3, while disk capacity increase to 177 PB (+25.36%). At the end of PQ11 (Jan 2013) the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to 347,307, which provide 3.86 Million HEP-SPEC 06.

The overall number of international and national VOs registered in the Operations Portal[1] at the end of March 2013 amounts to 212 (-6.2% from March 2012), including 22067 registered users (+5.36% increase from March 2012). The decrease in the number of VOs is due to a decommissioning campaign of inactive VOs that started in 2013.

The overall quantity of computing resources *used* in PY3 amounts to 12.01 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY2 amounted to 10.5 Billion HEP-SPEC 06 Hours). The PY3 workload was generated by 507.2 Million jobs, which amounts to an average of 1.43 Million job/day.

The overall compute resource utilization during PY3 has been significantly increasing both in terms of the cumulative number of jobs successfully done and the normalized CPU wall time consumed by all disciplines. In the refenrece period April 2012-March 2013 the rate of jobs succssfully executed incrased by +8.0%, while the total normalized CPU wall time (HEP-SEPC06) incrased by +45.8%.

While the HEP utilization is dominating in absolute terms (93.78% of the total EGI consumption in normalized CPU wall time hours), a number of other communities significantly increased their yearly CPU wall time utilization: Earth Sciences (+123.45%), Computational Chemistry (+78.31%), Astronomy Astro-particle and Astrophysics (+76.64%), Life Science (+65.12) and other sciences (+199.45%).

The performance of NGI services has been improving since January 2012 when the NGI Availability/Reliability statistics were introduced for the first time. Availability and Reliability were 98.17% and 99.98% respectively.

Monitoring of EGI.eu Core Infrastructure Platform was rolled to production in November 2012. The central EGI.eu services being monitored – which are part of the EGI Core Infrastructure Platform – are: the distributed monitoring infrastructure – SAM, the EGI-InSPIRE Metrics Portal, the Accounting Portal and central database, the central Operations Portal and the service registry GOCDB. The average availability and reliability performed by these EGI.eu tools in PQ11 is 98.6% and 98.6% respectively and exceeded the PY3 target (97%).

[1] http://operations-portal.egi.eu/vo

Various middleware stacks are in production in EGI. An indication of their distribution is given by the various Compute Element deployed by Resource Centres. CREAM-CE is in production in the 89.41% of the infrastructure, ARC-CE is second in deployment (0.11%) followed by GRAM (1.49%), Unicore6.TargetSystemFactory (1.49%) and QCG.Computing (1.12%).

gLite 3.1 and 3.2 software – released before the start of the EMI project[2] and still partially deployed by several RCs sites in PY3, is no longer supported. The decommissioning campaign of these two releases started in October 2012 and was successfully completed in PQ11. This first decommissioning campaign was subsequently followed by an EMI-1 decommissioning campaign which is still in progress (EMI-1 end of security updates and support is due on April 30 2013). The community of RCs participating to the early deployment of newly released software (Staged Rollout) has been expanding: the number of participating RCss has been progressively increasing to test a growing set of products from EMI, IGE and EGI-InSPIRE JRA1 (operational tools), and it currently amount to 74 teams.

SA1 contributed to the successful accomplishment of all EGI-InSPIRE objectives.

[2] http://www.eu-emi.eu/

TABLE OF CONTENTS

INDEX OF FIGURES

1 INTRODUCTION

This document illustrates the status and progress of the EGI production infrastructure and user community at the end of PY3. The current status of Resource Centres (RCs), Resource infrastructure Providers (RPs) and of the Operations Centres is show in Section 2. The amount of installed capacity being operated, the status and trends of the research communities supported by EGI, and their infrastructure usage are illustrated in Sections 3, 4 and 5 respectively. Section 6 illustrates the EGI service level management best-practices: it describes the current Operational Level Agreements in place, the agreed service level targets for services provided at the RC level, the RP level and EGI.eu level, and analyses the trends in performance delivered and the actions being implemented to support a number of infrastructures requiring consolidation. Section 7 describes the general core infrastructure services and the Cloud Infrastructure Platform is described in Section 8. The Staged Rollout infrastructure is described in Section 8. Section 10 provides information about Deployed Middleware Support Unit. Section 11 summarises the progress made in operations during PY3 to meet the project requirements and concludes the document.

2 RESOURCE INFRASTRUCTURE

This section provides information about the resource infrastructure of EGI encompassing Resource Centres (RCs), Resource infrastructure Providers (RPs) and the Operations Centres responsible of providing operational services to the community.

2.1 Resource Centres

A Resource Centre is the smallest resource administration domain in an e-Infrastructure. It can be either localised or geographically distributed. It provides a minimum set of local or remote IT Services compliant to well-defined IT Capabilities necessary to make resources accessible to Users. Access is granted by exposing common interfaces to Users [GLO].

Table 1. Number of EGI Resource Centres (March 2013).

Resource Centres	Number of RCs (certified)
EGI-InSPIRE Partners and NGI Council Members/PY3 Target	309 (does not include 31 suspended RCs)
From non-European EGI-InSPIRE Partners	28
From integrated Infrastructures (Canada, Latin America)	31
Total	**340/350**

As shown in the table above, the total number of certified RCs in March 2013 amounts to 340 instances, of which: 309 are contributed by European NGIs/EIROs that are EGI-InSPIRE partners or Council members and 31 by integrated RPs namely: Canada, Latin America – Brazil, Chile, Argentina, Venezuela and Mexico.

Of the 309 RCs mentioned above, 35 are contributed by Asia Pacific NGIs.

Figure 1 RCs distribution in March 2013 (data source: GOCDB and GSTAT).

The project target for PY3 of 350 RCs was not met (340 RCs certified centres against 350 expected), scoring a -4% deviation from the target. In PY3 the decrease was caused by the combination of different factors.

Firstly, the number of temporarily suspended RCs increased due to two grid middleware upgrade campaigns that were undertaken starting in PQ10[3]: one aimed at the replacement of unsupported gLite software, while the second was targeted to the decommissioning of EMI 1 software in preparation to its end of support in April 2013. Those RCs that had to be suspended as they could not adhere to the agreed retirement calendars will be integrated back into the production infrastructure as soon the infrastructure upgrade is completed.

A second factor that contributed to this reduction was the infrastructure consolidation campaign undertaken by the Italian NGI at the beginning of PY03: various Italian RCs were decommissioned in order to merge resources into bigger sites and take out of production those which were not providing sufficient availability. In the same quarter, additional RCs were suspended mainly in Armenia, Belarus, Brazil and Bulgaria because of Availability and Reliability issues.

Finally, because of sustainability issues, both the Irish Operations Centre and IGALC (Iniciativa de Grid de America Latina – Caribe, formerly supported by the GISELA EC project) stopped their operations at the end of December 2012. Irish RCs were all closed, while part of the IGALC RCs migrated to the Latin America federated Operations Centre. The loss of production resources in Ireland and Latin America was however compensated by an extension of the infrastructure in the Asia Pacific region. In March 2012 the total number of certified RCs decreased from 352 to 291, it then increased to 306 (end of PQ8), to 310 at the end of PQ9, 326 in PQ10) and 332 to date. Figure 2 plots the number of RCs from January 2009.

The decrease in the number of production RCs was however compensated by a substantial increase in the offered capacity: compute resources increased by +33.6% in PY3, while disk capacity increase to 177 PB (+25.36%).

In PY2 two new RPs got engaged with EGI by signing a Resource infrastructure Provider MoU: the South African Grid Initiative [SAG] and the Ukrainian National Grid [UNG]. In PY3 the integration of the Ukrainian National Grid was successfully finalized – NGI_UA comprises 12 production RCs to date, while the integration of South African Grid Initiative was put on hold due to lack of local effort and will be likely resumed in PY4.

GOCDB[4] was used to extract information about the numbers of certified production RCs reported in this section.

[3] https://wiki.egi.eu/wiki/Software_Retirement_Calendar

[4] https://goc.egi.eu

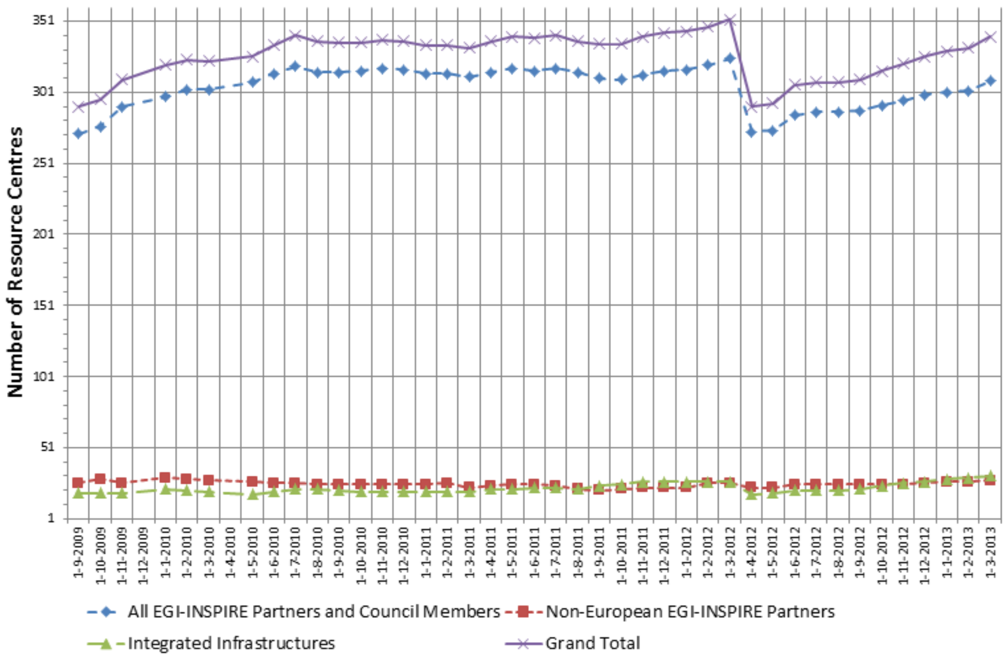

Figure 2. Number of certified production RCs from 01/09/2009 to 01/03/2013. The drastic reduction in number in April 2012 is due to both the decommissioning of small RCs in Italy and to suspensions in other NGIs because of low performance issues (data source: GOCDB).

2.2 Resource infrastructure Providers

The Resource infrastructure Provider is the legal organisation responsible for any matter that concerns their respective Resource Infrastructure. It provides, manages and operates (directly or indirectly) all the operational services required to an agreed level of quality as required by the Resource Centres and their user community. It holds the responsibility of integrating these operational services into EGI in order to enable uniform resource access and sharing for the benefit of their Users. The Resource infrastructure Provider liaises locally with the Resource Centre Operations Managers, and represents the Resource Centres externally. Examples of a Resource infrastructure Providers are the European Intergovernmental Research Organisations (EIRO) and the National Grid Initiatives (NGIs) [GLO].

In March 2013 EGI comprises resources provided across 55 countries and one European Intergovernmental Research Institute – CERN.

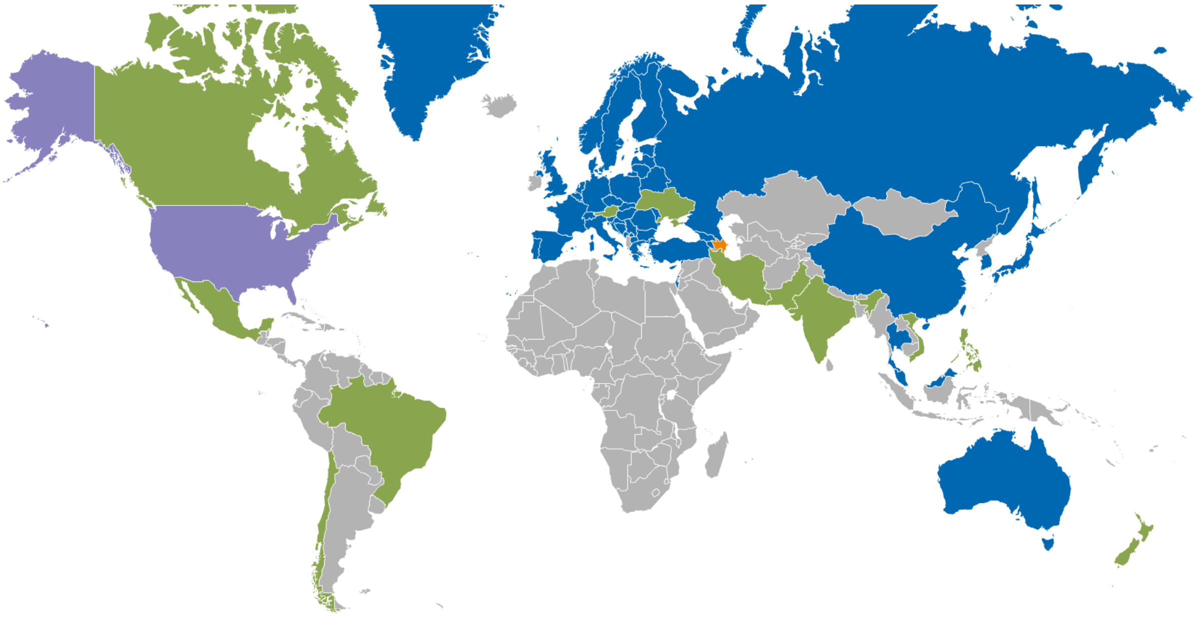

Figure 3 **RPs distribution in March 2013 (data source: GOCDB). Legend: (blue) Integrated EGI-InSPIRE Partners and EGI Council Members, (green) External Resource Providers, (orange) Internal/External Resource Providers – Azerbaijan, (purple) Peer Resource Providers – Open Science Grid.**

EGI-InSPIRE partners or RPs that are Council members contribute resources from 44 countries. However, 11 countries contribute resources through Resource infrastructure Providers that are non-EGI-InSPIRE partners but are fully integrated with the EGI Services Infrastructure. These are:

- India, Pakistan, Vietnam and New Zealand (Asia Pacific Federation);
- Austria (Italian Federation): operations support of Austrian RCs moved from NGI_NDGF to NGI_IT during PY2;
- Canada (Canada Federation);
- Brazil, Chile, Argentina, Venezuela and Mexico (Latin America Federation).

The distribution of RCs per country and per Operations Centres is reported in Table 2.

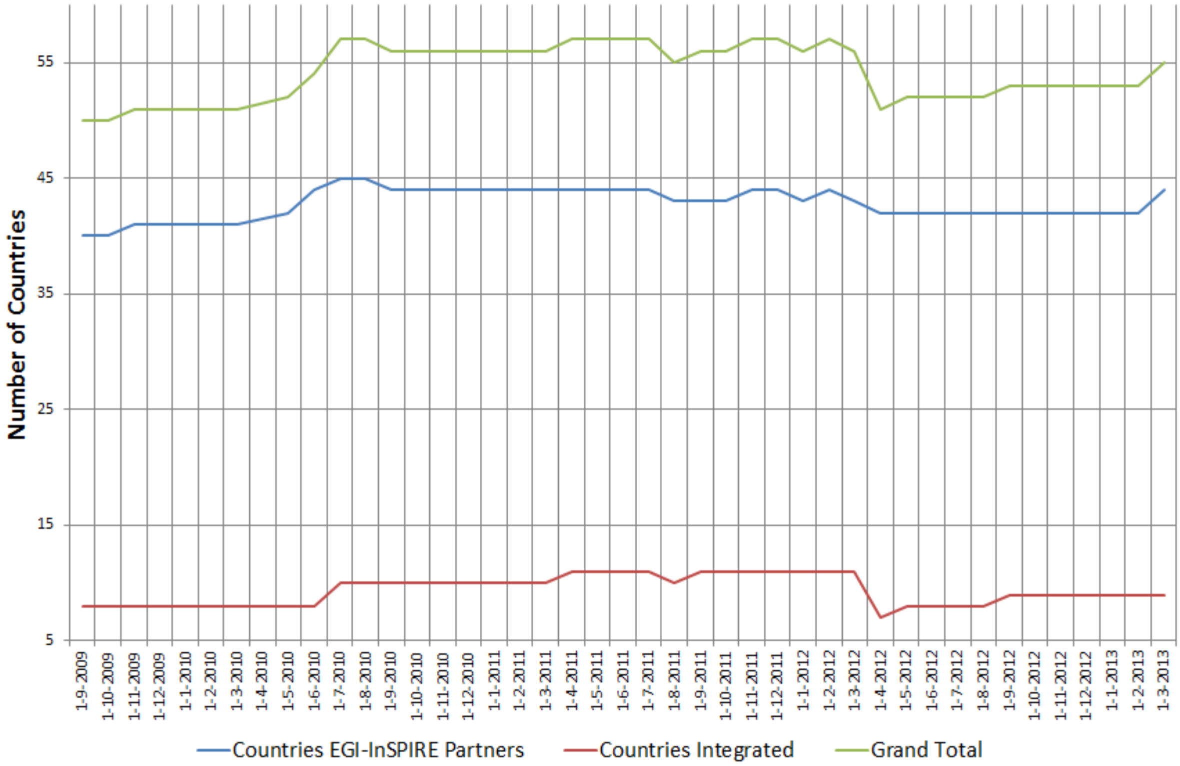

Figure 4. EGI countries hosting certified production Resource Centres from 01/09/2009 to 01/03/2013 (data source: GOCDB).

The number of countries contributing resources is approximately constant. In PY3 two new countries started contributing resources: Iran and Vietnam.

Since PQ8 the production infrastructure in Armenia has been unstable, and RCs have been periodically temporarily removed from the infrastructure; the infrastructure reached stability in PQ11 with two production RCs. The growth trend since 01 September 2009 is illustrated in Figure 4. In March 2013, the EGI-InSPIRE partners that are still not contributing resources are: Albania, Argentina, Indonesia and Singapore.

Table 2. Distribution of production certified RCs across countries and Operations Centres in PQ8, PQ9, PQ10, PQ11 and March 2013 (data is sorted by country). (Data source: GOCDB).
I = Integrated, EP = European Partner, P = non-European partner, C = only Council member

Operations Centre	Country		PQ8	PQ9	PQ10	PQ11	01/03/2013
NGI_ARMGRID	Armenia	EP	0	0	0	0	2
Asia Pacific	Australia	P	1	1	1	1	1
NGI_IT	Austria	I	2	2	2	2	2
NGI_BY	Belarus	EP	1	1	1	1	1
NIG_NL	Belgium	C	3	3	3	3	3
NGI_BA	Bosnia and H.	EP	1	1	1	1	1
Latin America	Brazil	I	2	2	2	2	2
NGI_BG	Bulgaria	EP	4	4	4	4	4
Canada	Canada	I	7	7	7	8	9
Latin America	Chile	I	1	1	1	1	1
Asia Pacific/Canada	China	P	3	3	3	3	3
NIG_HR	Croatia	EP	3	3	3	3	3
NGI_CY	Cyprus	EP	1	1	1	1	1
NGI_CZ	Czech Republic	EP	2	2	2	2	2
NGI_NDGF	Denmark	EP	2	2	2	2	2
NGI_NDGF	Estonia	C	2	2	2	2	2
NGI_FI	Finland	EP	2	11	11	11	11
NGI_FRANCE	France	EP	18	18	18	18	18
NGI_MARGI	FYROM	EP	3	3	3	3	3
NGI_GE	Georgia	EP	1	1	1	1	1
NGI_DE	Germany	EP	18	18	19	19	19
NGI_GRNET	Greece	EP	14	14	15	16	16
NGI_HU	Hungary	EP	3	3	3	3	3
Asia Pacific	India	I	0	1	2	2	2
Asia Pacific	Iran	I	0	0	1	1	1
NGI_IE	Ireland	EP	6	0	0	0	0
NGI_IL	Israel	EP	5	5	5	6	7
NGI_IT	Italy	EP	49	50	50	53	56
Asia Pacific	Japan	P	3	3	3	3	3
NGI_NDGF	Latvia	EP	2	2	2	2	2
NGI_NDGF	Lithuania	EP	1	1	1	1	1
Asia Pacific	Malaysia	P	4	5	5	5	5
Latin America	Mexico	I	1	1	1	1	1
NGI_ME	Montenegro	EP	1	1	1	1	1
NGI_NL	Netherlands	EP	13	15	15	16	16
Asia Pacific	New Zealand	I	1	1	1	1	1
NGI_NDGF	Norway	EP	1	1	1	1	1
Asia Pacific	Pakistan	P	1	1	1	1	1
Asia Pacific	Philippines	P	1	1	1	1	1
NGI_PL	Poland	EP	10	10	10	10	10
NGI_IberGrid	Portugal	EP	5	5	6	6	6
NGI_RO	Romania	EP	10	10	10	10	10
Russia	Russia	EP	10	10	10	10	11
NGI_AEGIS	Serbia	EP	5	5	5	5	5
NGI_SK	Slovakia	EP	4	4	4	4	4
NGI_SI	Slovenia	EP	2	2	2	2	2
Asia Pacific	South Korea	P	4	4	4	5	5
NGI_IberGrid	Spain	EP	21	21	21	21	21
NGI_NDGF	Sweden	EP	1	1	1	1	1
NGI_CH/CERN	Switzerland	EP	7	7	7	8	8
Asia Pacific	Taiwan	P	4	4	4	5	5
Asia Pacific	Thailand	P	2	3	3	3	3
NGI_TR	Turkey	EP	3	3	3	3	3
NGI_UA	Ukraine	I	4	6	7	11	12
NGI_UK	United Kingdom	EP	22	22	24	24	24
Asia Pacific	Vietnam	P	0	0	0	0	1
TOT Resource Centres			297	308	316	330	340
TOT Countries			52	52	53	53	55

2.3 Operations Centres

The Operations Centre is defined to be a centre offering operations services on behalf of the Resource infrastructure Provider [GLO], and it can serve multiple RPs.

EGI currently comprises 28 national operations centres and 7 federated operations centres encompassing multiple NGIs (Table 3). The existing federated centres in Europe (IberGrid, NGI_NL and NGI_IT) each contain two countries and are the result of a collaboration agreement that is expected to continue in PY4. In contrast, integrated federated centres in Asia Pacific and Latin America encompass a large number of countries. The creation of new national grid initiatives in those regions will depend on their expansion plans and on national policies.

A new operations centre was created in Ukraine (NGI_UA) in June 2012. Ukrainian resources were previously operated by ROC Russia.

Two Operations Centres were decommissioned: ROC IGALC and NGI IE. Most of the ROC IGALC RCs were handed off to ROC LA.

Federated operations centres	Member countries	Comments
Asia Pacific	Australia, China, India, Japan, Malaysia, New Zeeland, Pakistan, South Korea, Taiwan, Thailand, Vietnam, Philippines	
Canada	Canada, China	
IberGrid	Portugal, Spain	
Latin America (ROC_LA)	Brazil, Chile, Colombia, Mexico, Argentina, Brazil, Chile, Venezuela,	Part of the RCs from Iniciativa de Grid de America Latina – Caribe (ROC_IGALC) were moved to ROC_LA and ROC_IGALC was decommissioned
Italy (NGI_IT)	Austria, Italia	Operations of Austrian sites provided by NGI_IT since November 2011
Nordic countries and Baltic region (NGI_NDGF)	Denmark, Estonia, Latvia, Lithuania, Norway, Sweden	
Netherlands (NGI_NL)	Belgium, Netherlands	

Table 3. List of EGI federated Operations Centres

3 INSTALLED CAPACITY

Installed capacity is monitored at the end of each project quarter (PQ). Metrics are automatically collected from the Information Discovery System and validated by NGIs, consequently statistics herein reported depend on the accuracy of information published and on the responsiveness of NGIs in reporting resources that are not publishing.

ARC, UNICORE (distributed through EMI releases) and GLOBUS services are currently not published in the Information Discovery System with the exception of ARC resources operated by NGI_NDGF. Because of this the statistics below do not provide information about ARC, GLOBUS and UNICORE RCs. The EMI release v.3 and IGE release v. 3.1 allow publishing of information into top-BDII. A top-BDII and site-BDII deployment plan for countries offering ARC, GLOBUS and UNICORE is part of the PY4 activity plan.

3.1 Compute Resources

At the end of PQ11 the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to 347,307, which provide 3.32 Million HEP-SPEC 06, while the total number including compute resources contributed by integrated and peer infrastructures amounts to 373,235 units. This value significantly exceeds the PY3 target of 300,000 total cores.

Looking at the compute resources provided by EGI partners (EGI-InSPIRE partners and EGI Council members) – Table 5, the number of CPU cores increased by 33.6% since March 2012, while the installed capacity in HEP-SPEC 06 increased by 44.7%. This increase reflects the advancements in CPU technology and core density per CPU and shows a trend in building up capacity in the existing RCs.

Table 4. EGI-InSPIRE logical CPUs

Logical CPUs	PQ11/PY3 Target
EGI Council participants	347,307/300,000
EGI-InSPIRE partners plus integrated and peer infrastructures	373,235

Table 5. Installed compute capacity in EGI-InSPIRE partners and EGI Council members (logical CPUs and Million HEP-SPEC 06) in April 2010, March 2011, March 2012 and March 2013

	April 2010 (EGEE-III Infrastructure)	March 2011 (EGI-InSPIRE Infrastructure)	March 2012 (EGI-InSPIRE Infrastructure)	March 2013 (EGI-InSPIRE Infrastructure)
Logical CPUs/yearly increase	192,000	207,203/+7.9%	270,800/+30.7%	373,235/+33.6%
Million HEP-SPEC 06	1.34	1.98/+47.7%	2.96/+49.5%	3,86/+44.7%

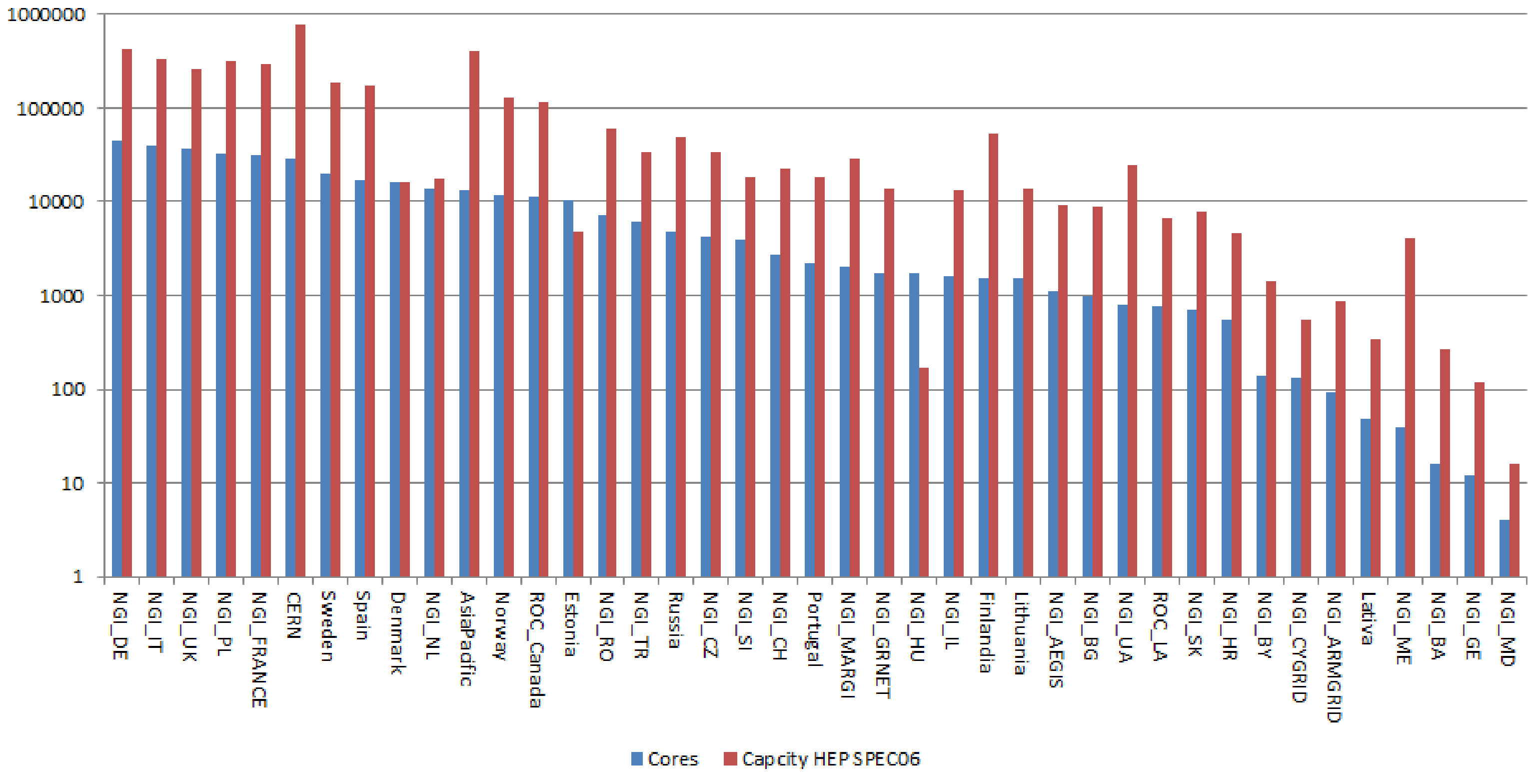

Figure 5. Log scale distribution of logical cores (blue bar) and HEP SPEC 06 installed capacity (red bar) at the end of PQ11 across EGI Resource infrastructure Providers, including EGI-InSPIRE partners, EGI Council members and integrated infrastructures. Data sorted by number of cores. Source: project quarterly metrics and top-BDII.

3.2 Storage Resources

Information from each RC about the storage capacity is periodically collected by the Metrics Portal from the Information Discovery System (Gstat) and validated by EGI-InSPIRE partners. As for compute capacity, the accuracy of data available from the Information Discovery System depends on the availability of correct and up to date information as provided by the storage dynamic information providers installed at RCs.

At the end of PQ11, the total amount of reported installed disk capacity amounts to 177 PB (+25.36% yearly increase). The distribution of disk storage resources among the EGI-InSPIRE partners is illustrated in Figure 6, which shows that disk capacity is concentrated across six NGIs/EIROs, which are in descending order (for PQ11): Germany, CERN, United Kingdom, Italy, France and Spain.

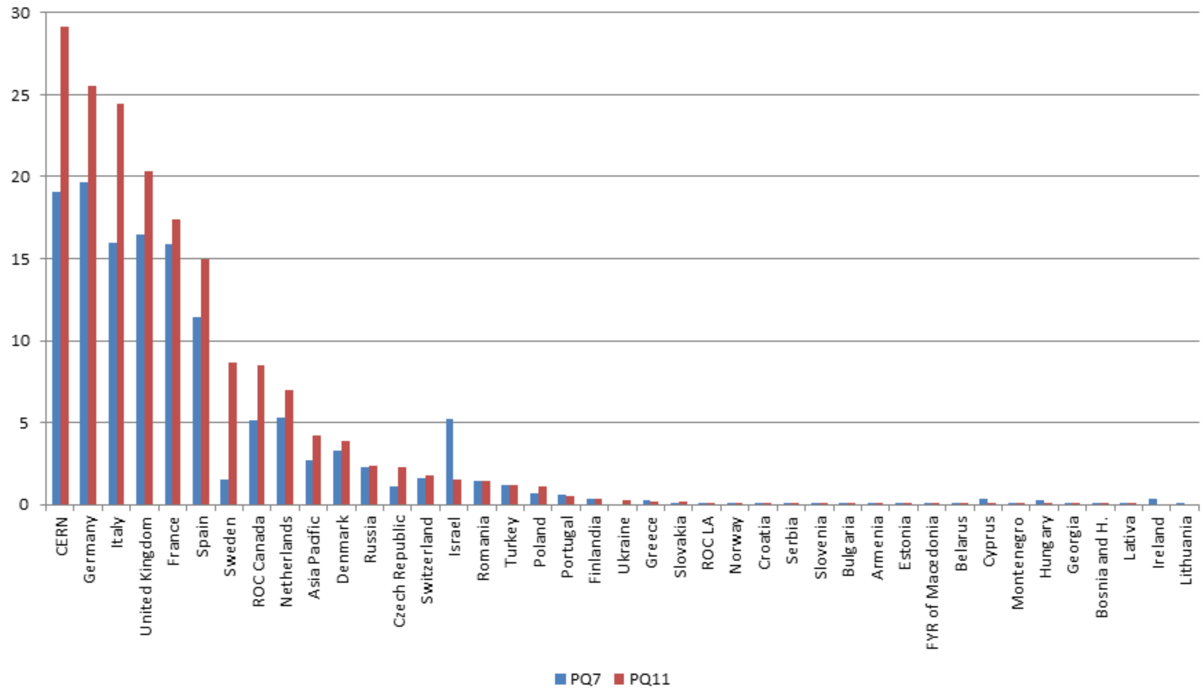

Figure 6. Installed disk capacity in PB across the EGI RPs at the end of PQ11 – red bar – compared to the installed capacity in PQ7 – blue bar (source: Metrics Portal and Gstat).

Tape capacity is mainly provided by CERN and WLCG Tier-1 RCs. At the end of PQ11 the total installed tape (also known as *nearline*) capacity reported in Gstat amounts to 187.6 PB (+28% yearly increase).

3.3 Compute Resources for Parallel Jobs

Information about the number of high-performance clusters operated is gathered periodically in the project quarterly reports. With high-performance we refer to clusters that feature a local high-speed low-latency interconnect (e.g. Myrinet[5], InfiniBand[6]). The clusters that qualify as high-performance,

[5] http://www.myri.com/myrinet/overview/

[6] http://www.infinibandta.org/

as reported by the Resource infrastructure Providers, amount in total to 42 units (4.8% yearly increase with respect to PQ7) at the end of PQ11.

Table 6. Integration metrics (HPC and MPI)

Metric	PQ7	PQ11/PY3 Target
Number of HPC clusters (M.SA1.Integration.1)	40	42/50
Number of sites with MPI (M.SA1.Integration.2)	108	89/120

At the end of PQ11 Message Passing Interface [MPI] jobs were supported by 89 sites (-17.6% yearly decrease) as shown in Table 6. The steady increase observed in PY1 and PY2 was unexpectedly followed by a change of trend in PY3. The reason for this needs further investigation.

As a result of the works of the MPI Virtual Team[7], starting with PQ12 a new mechanism for registering and monitoring resources offering the MPI capability will be rolled out to production: information about MPI capabilities will not only be published by services to the Information Discovery Service as today, but it will also be registered into the EGI service registration facility (GOCDB). A more accurate mechanism to estimate the MPI support in the infrastructure will be available. In addition, starting with the EMI 3 release, the new APEL publisher will be capable of reporting accounting information of multi-core jobs. The new publisher will be gradually deployed by the production infrastructure. Accounting of MPI jobs will be a more accurate indicator of the amount of parallel computing workload supported by EGI.

During PY3, HPC integration activities focused on the objective of supporting loosely and tightly coupled multi-scale simulations across EGI and PRACE resources continued. Pilot activities were carried out in collaboration with the MAPPER project[8], and a new initiative seeing the involvement of EGI, EUDAT[9], PRACE and user communities started in November 2012[10] aiming for the integration of data access and processing across the three infrastructures. Objectives of this initiative are[11]:

- To collect use cases for data access, transfer, replication and processing among different e-Infrastructures (EGI, EUDAT and PRACE): technological and operational barriers shall be removed or mitigated as far as possible.
- To identify common data access and transfer tools and protocols that can be provided by all three e-infrastructures, and that are useful for the collaborating user communities.
- To identify technology, operational and/or organisational gaps and suggest improvements if use cases cannot be realized across the three e-infrastructures.

[7] MPI Virtual Team: https://wiki.egi.eu/wiki/VT_MPI_within_EGI

[8] www.mapper-project.eu/

[9] http://www.eudat.eu/

[10] EGI, EUDAT and PRACE workshop on data management:

https://indico.egi.eu/indico/conferenceTimeTable.py?confId=1228#20121126

[11] https://confluence.csc.fi/pages/viewpage.action?pageId=28837071

4 DISCIPLINES, VIRTUAL ORGANIZATIONS AND USERS

This section provides information about the evolution of the user community (users registered in VOs) in some of the main scientific disciplines currently identified by EGI at the infrastructure level, namely: Computer Science and Mathematics, Multidisciplinary VOs, Astronomy Astrophysics and Astro-Particle Physics, Life Sciences, Computations Chemistry, Earth Sciences, Fusion, High-Energy Physics, Infrastructure, and Others[12]. A new science classification was proposed in PY3 and discussed with user communities, which will adopted in PY4.

The overall number of international and national VOs registered in the Operations Portal[13] at the end of March 2013 amounts to 212 (-6.2% from March 2012), including 22067 registered users (+5.36% increase from March 2012). The decrease in the number of VOs is due to a decommissioning campaign that started in 2013 targeted at inactive VOs.

Currently user statistics extracted from VO Membership Services do not provide information about the number of active users that by interacting with the grid through scientific gateways, are associated to robot certificates.

4.1 VO Distribution across scientific fields

The disciplines which recorded the larger number of VOs during PY3 are: High-Energy Physics (-7.7% yearly *relative* decrease, with +3.92% of new users registered yearly), Others (-10.8%, with -4.71% of new users), Multidisciplinary VOs (-7.9% VOs and +14.03% users), Infrastructure (-3.4% VOs and +6.6% users), Astronomy Astrophysics and Astro-particle Physics (constant number of VOs and -74.57% users), Life Sciences (+5.6% VOs and +0.85% users), Earth Sciences (-16.7% VOs and +2.78% users), Computer Science and Mathematics (-16.7% VOs and +11.90% users), Computational Chemistry (constant number of VOs and +1.65% users) and Fusion (constant number of VOs and +4.17% users).

The distribution of VOs per discipline is illustrated in Figure 8. Disciplines that in PY3 have increased their fraction of users in comparison with the total number are: Infrastructure (13.2%), Astronomy Astrophysics and Astro-particles Physics (12.7%) and Life Sciences (9%).

12 "Others" is a category of user communities that do not belong to the other disciplines that are part of the current classification. The scientific discipline classification of EGI is being reviewed.

13 http://operations-portal.egi.eu/vo

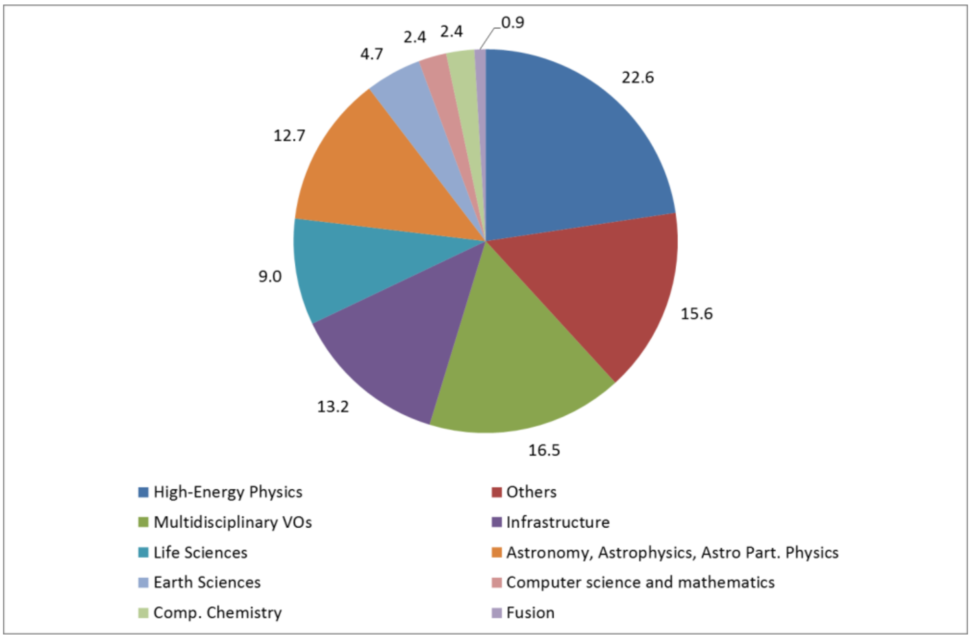

Figure 7. Distribution of number VOs per discipline (March 2013, source: Operations Portal).

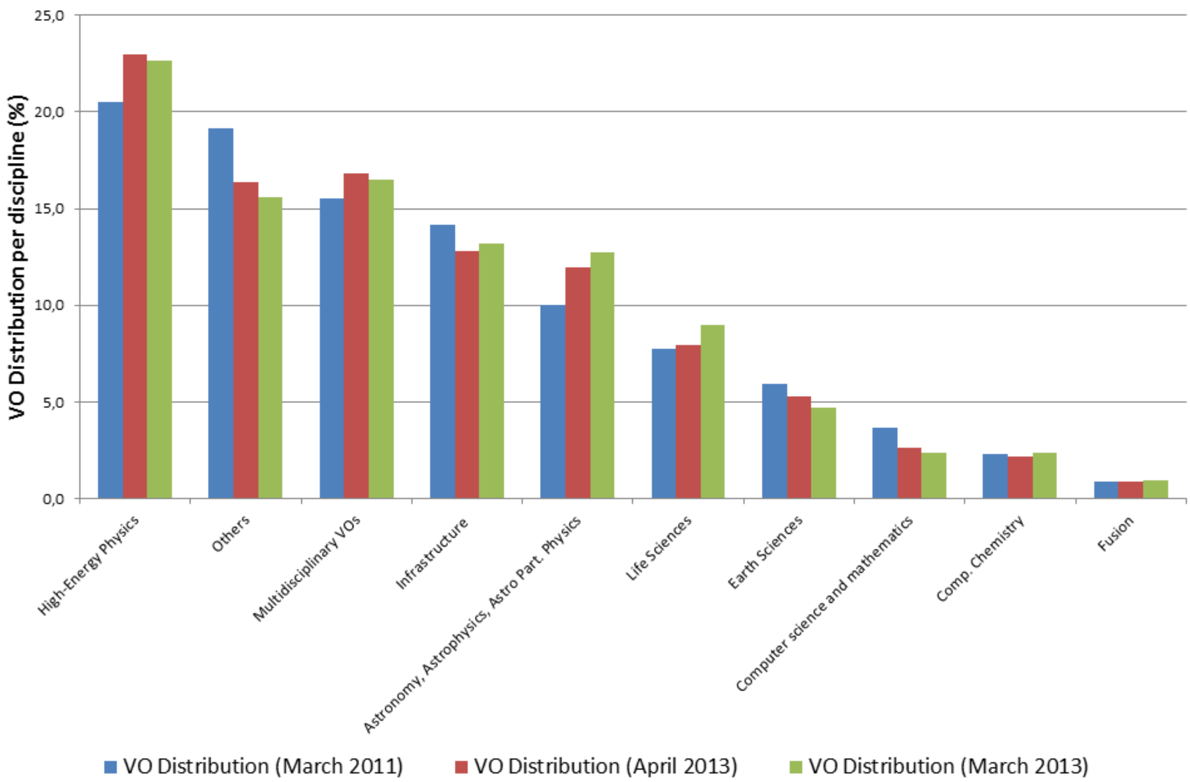

Figure 8. Comparison of the VO distribution at the end of March 2011 (blue bars), at the end of March 2012 (red bars) and at the end of March 2013 (green bars). Source: Operations Portal.

4.2 User Distribution across scientific fields

The largest disciplines in terms of number of registered users are: High Energy Physics (38.60%), Others (17.50%) and Multidisciplinary VOs (16.47%). During PY3 the number of registered users for some disciplines has increased: Infrastructure (+6.6%), Multidisciplinary VOs (+13.03%), Astronomy Astrophysics and Astro-particle Physics (+6.81) and Computer Science (+11.90%). The detailed user distribution per discipline is presented in Appendix.

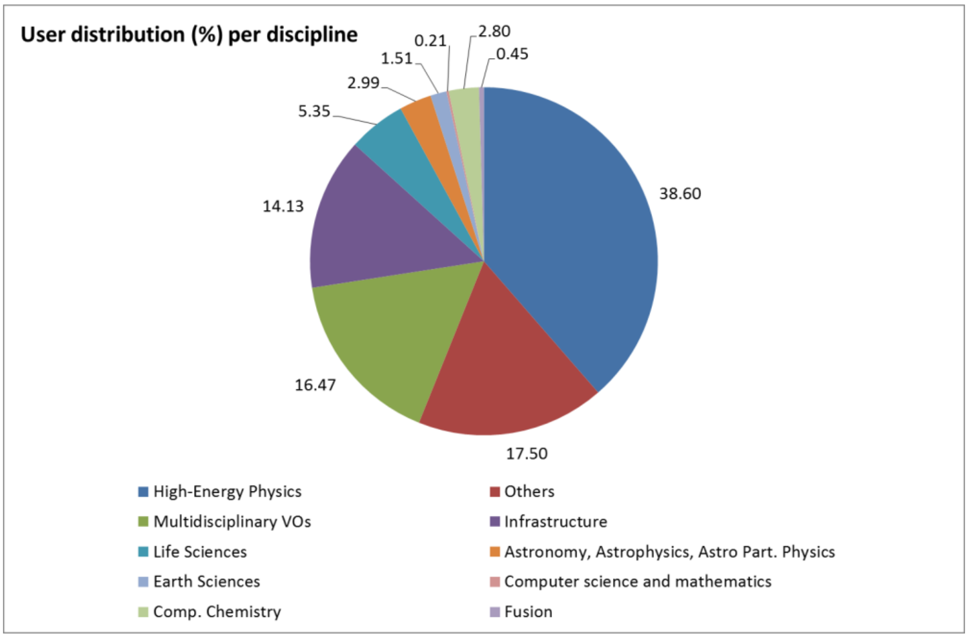

Figure 9. User distribution per discipline (March 2013, source: Operations Portal)

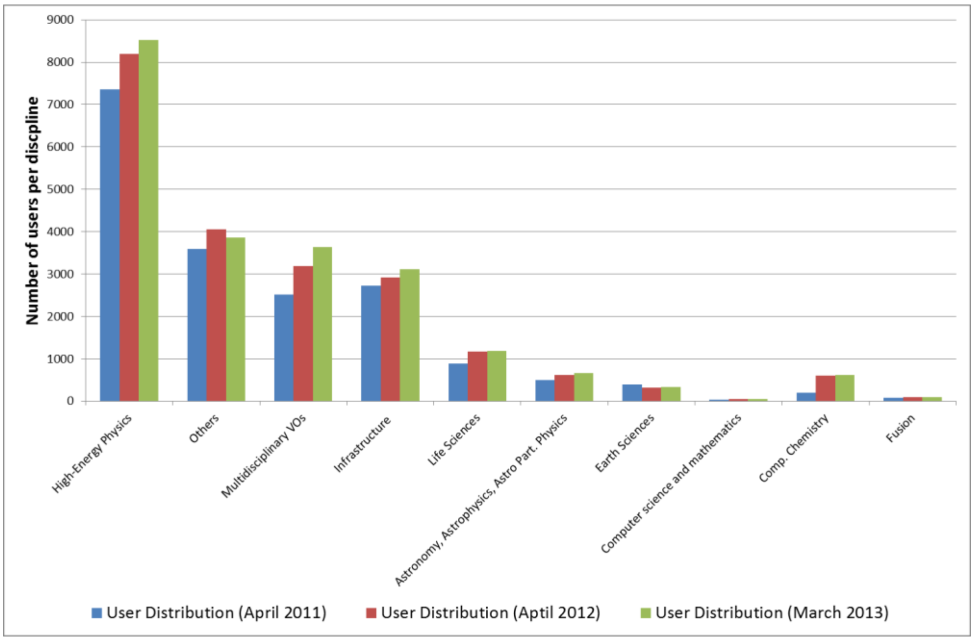

Figure 10. Comparison of the number of users per discipline in April 2011 (blue bar), April 2012 (red bar) and at the end of March 2013 (green bars). Source: Operations Portal.

4.3 Resource Utilization per Discipline

Table 7. Cumulative compute resource utilization (number of executed jobs and normalized CPU wall time). Source: Accounting Portal

Metric	Period	Metric value	Yearly increase (%)
Cumulative number of executed jobs - Million	Apr 2010 - Mar 2011	320.7	
	Apr 2011 – Mar 2012	485.1	+51.3%
	Apr 2012 – Mar 2013	524.0	+8.0%
Normalized CPU wall time (HEP-SPEC06) – Billion hours	Apr 2010 - Mar 2011	6.37	
	Apr 2011 – Mar 2012	10.27	61.3%
	Apr 2012 – Mar 2013	14.99	45.8%

The overall compute resource usage during PY3 has increased significantly both in terms of the cumulative number of jobs successfully done and the normalized CPU wall time consumed by all disciplines. In the reference period April 2012-March 2013 the rate of jobs succssfully executed incrased by +8.0%, while the total normalized CPU wall time (HEP-SPEC06) incrased by +45.8%. Table 7 compares the April 2012 – March 2013 increase trends with those achived in the previous 2 years.

HEP-SPEC 06 is the EGI reference performance benchmark of compute resources [HS06]. It was defined by the HEPiX Benchmarking Working Group and it is based on SPEC. One HEP-SPEC 06 corresponds approximately to 250 SI00 (this was tested with HEP applications).

As the CPU performance varies greatly between different resources, even within a single site, a reference is needed to provide a fair comparison of resource usage consumption. The APEL accounting system used in EGI scales CPU time to a reference benchmark of 1,000 SI2K hours (4 HEP-SPEC 06 hours). Each Grid site publishes a value for the CPU speed (described by the SpecInt2000 performance benchmark) for each site cluster as part of the site's GLUE schema. When generating accounting records, APEL queries the site's Information Discovery System to obtain this data. Each individual record will then contain the CPU speed equivalent from the worker node where the job was executed. Once the record has been published into the APEL Accounting Server, the CPU time can then be normalized to the reference value (4 HEP-SPEC 06 hours).

The trend in usage of normalized CPU wall clock since the beginning for EGI-InSPIRE, during PY2 and PY3 are shown in Figure 11(a).

The High-Energy Physics discipline (contributing 38.60% of the user community) utilizes the highest amount of resources: 93.78% of the overall EGI amount of normalized CPU wall time hours consumed. As indicated in Table 8, the HEP usage yearly increase amounts to 40.97%.

While the HEP utilization is dominating in absolute terms, a number of other communities significantly increased their CPU wall time utilization: Earth Sciences (+123.45% yearly increase), Computational Chemistry (+78.31%), Astronomy Astro-particle and Astrophysics (+76.64%), Life Science (+65.12) and other sciences (+199.45%).

Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.82% of the overall EGI used CPU wall clock time. Life Sciences are the third community for usage (1.52% of the overall EGI used normalized CPU time). For this community the job submission pattern shows a further increase of the job submission rate (+156.79%). This trend already started in PY2, during which the rate already showed a relative increase of +42.54%.

Computational Chemistry, Earth Science also increased both their used CPU wall clock time and the job rate. The overall trend of used normalized CPU wall clock time for non-HEP disciplines is plotted in Figure 13. As the diagram shows, for the largest user communities with the exception of HEP, usage is subject to short-term fluctuations, and the job workload produced is generally independent from the corresponding amount of CPU wall time consumed.

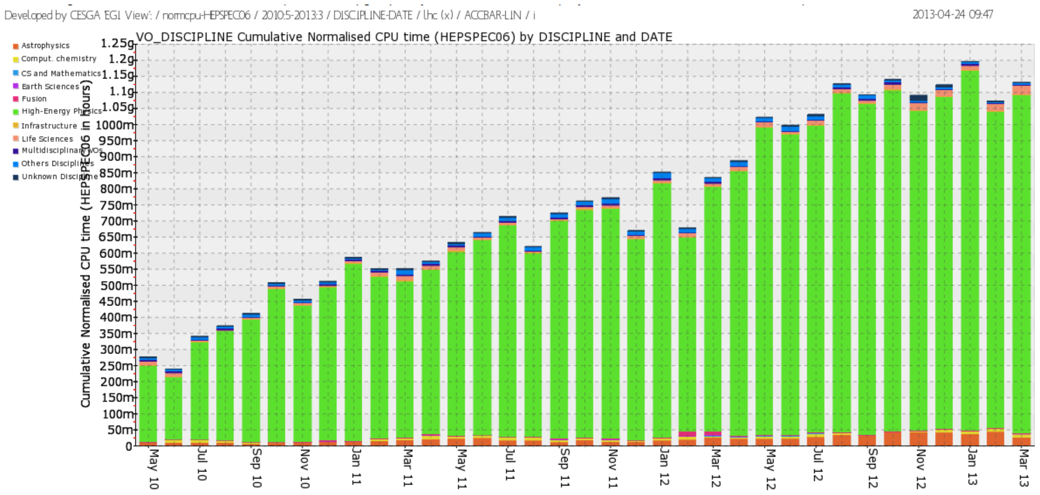

(a)

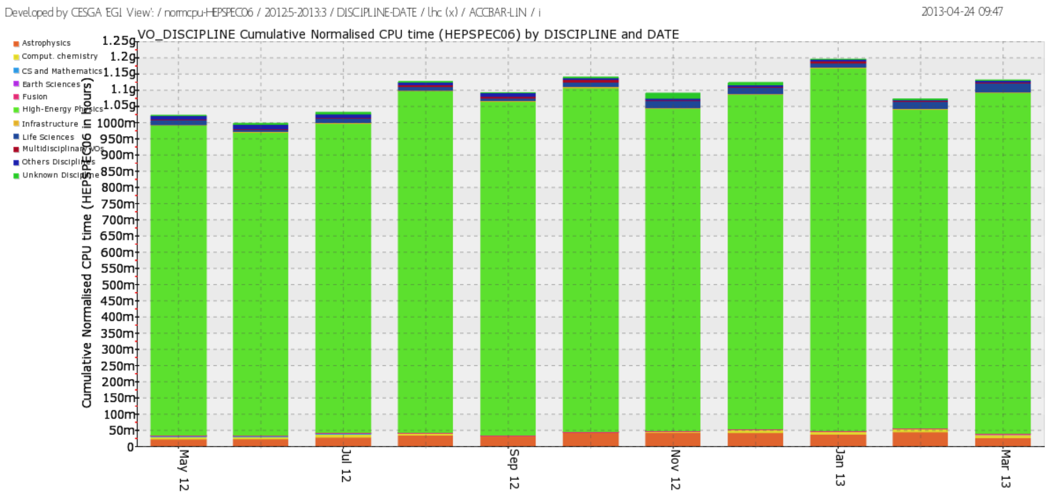

(b)

Figure 11. Usage of EGI resources (HEP-SPEC 06 CPU wall clock hours) from (a) the beginning of the project to date, and (b) during PY3 (source: accounting portal).

Table 8. Used normalized CPU wall clock time and jobs done per discipline in PY2, PY3 and the respective yearly increase (source: accounting portal).

Discipline	May 12–April 13		May 11 – April 12		Jobs	CPU wall time
	% CPU n. wall time (A)	% of Jobs done (B)	% CPU n. wall time (C)	% of jobs done (D)	(yearly increase from May 11) (E)	(yearly increase from May 11) (F)
High-Energy Physics	93.78	89.58	93.60	91.58	**+1.22%**	**+40.97%**
Infrastructure	0.10	2.88	0.20	3.26	-8.70%	-29,67%
Life Sciences	1.52	4.34	1.30	1.75	**+156.79%**	**+65.12%**
Astrophysics	2.82	1.82	2.25	1.58	**+18.57%**	**+76.64%**
Multidisciplinary	0.12	0.17	0.39	0.48	-62.77%	-56.97%
Others Disciplines	0.59	0.45	1.23	0.72	-36.713%	-32.12%
Unknown Discipline	0.43	0.27	0.20	0.29	**-3.08%**	**+199.45%**
Comput. Chemistry	0.48	0.22	0.38	0.03	**+83.04%**	**+78.31%**
Fusion	0.01	0.10	0.37	0.13	-24.56%	-96,98%
Earth Sciences	0.15	0.11	0.10	0.05	**+139.95%**	**+123.45%**
CS and Mathematics	0.00	0.07	0.00	0.03	+170.56%	-68.06%

Figure 12. Distribution of consumed normalized CPU wall time among the main active VOs (May 2012-March 2013)Figure 12 shows the distribution of consumed normalized CPU wall time among the main VOs.

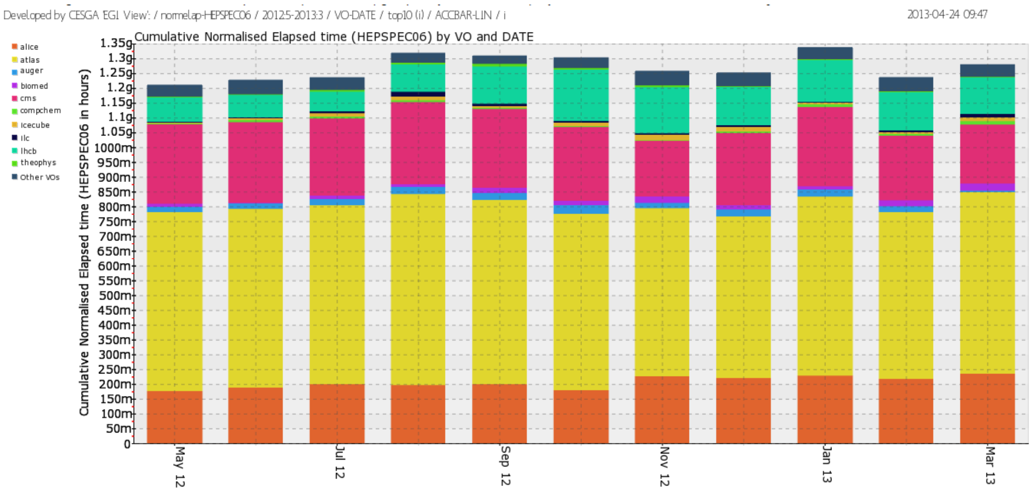

Figure 12. Distribution of consumed normalized CPU wall time among the main active VOs (May 2012-March 2013)

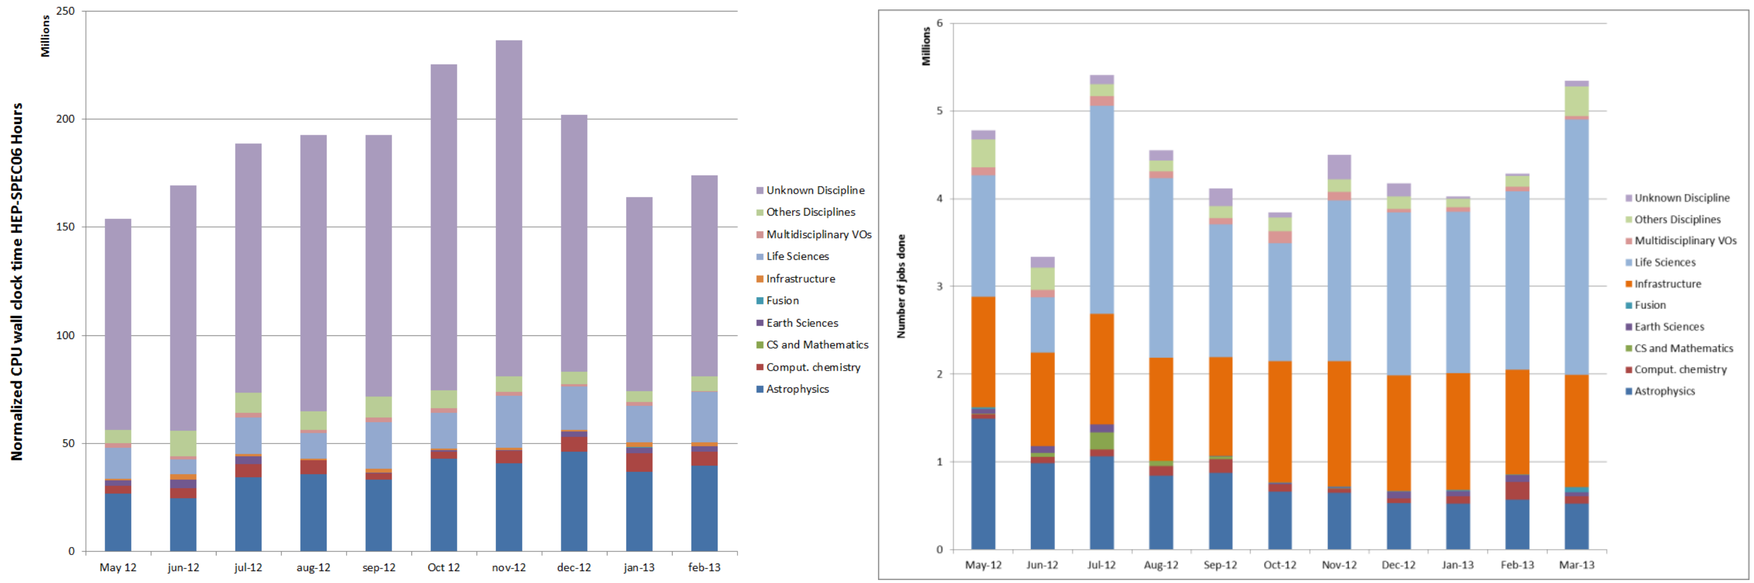

Figure 13. Used normalized CPU wall clock time (left) and number of jobs done (right) across disciplines during PY3.

5 RESOURCE USAGE

EGI accounting information is gathered and stored centrally for display through the accounting portal[14]. Accounting information is aggregated by Operations Centre, whose list is obtained from GOCDB.

Table 9. Annual compute resource usage (yearly figures)

	PY2	PY3 Value/Target
Total normalized CPU wall clock time consumed (Billion HEP-SPEC 06 hours)	10.5	12.01
Jobs per year (Million)	492.5	507.2
Average number of Jobs per day (Million)	1.35	1.43/1.2

The overall quantity of computing resources used in PY3 amounts to 12.01 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY2 amounted to 10.5 Billion HEP-SPEC 06 Hours) as shown in Table 9. The PY3 workload was generated by 507.2 Million jobs, which amounts to an average of 1.43 Million job/day.

PY3 usage expressed in HEP-SPEC 06 Hours of CPU wall time across the various resource infrastructures of EGI is plotted in Figure 14, where infrastructures are grouped by operations centre. The diagram also shows the distribution between HEP user communities (blue bars) and the non-HEP user communities (green bars), the top infrastructures for multidisciplinary support being (in decreasing order): NGI_UK, NGI_DE, NGI_IT, NGI_FR and CERN. Usage distribution naturally reflects availability of installed capacity (Section 3), however the level of multidisciplinary support varies considerably across the infrastructures. Figure 15 plots the distribution of used HEP-SPEC 06 CPU wall clock hours of non-HEP user communities. NGI_DE is the infrastructure with the largest absolute amount of resources used by non-HEP communities with more than 203 Million CPU wall time hours, followed by NGI_IT, NGI_FR, NGI_NL and NGI_UK.

The Figure 15 shows how support of HEP is dominant in large resource infrastructures, while other disciplines dominate in various countries in Eastern Europe, where is some cases it equals 100% of the entire usage of resources. An expected outcome of the EGI-InSPIRE activities in outreach and technical support of new user communities introduced in PY3, is that the fraction of non-HEP usage will increase in future years.

[14] http://www4.egee.cesga.es/accounting/egee_view.php

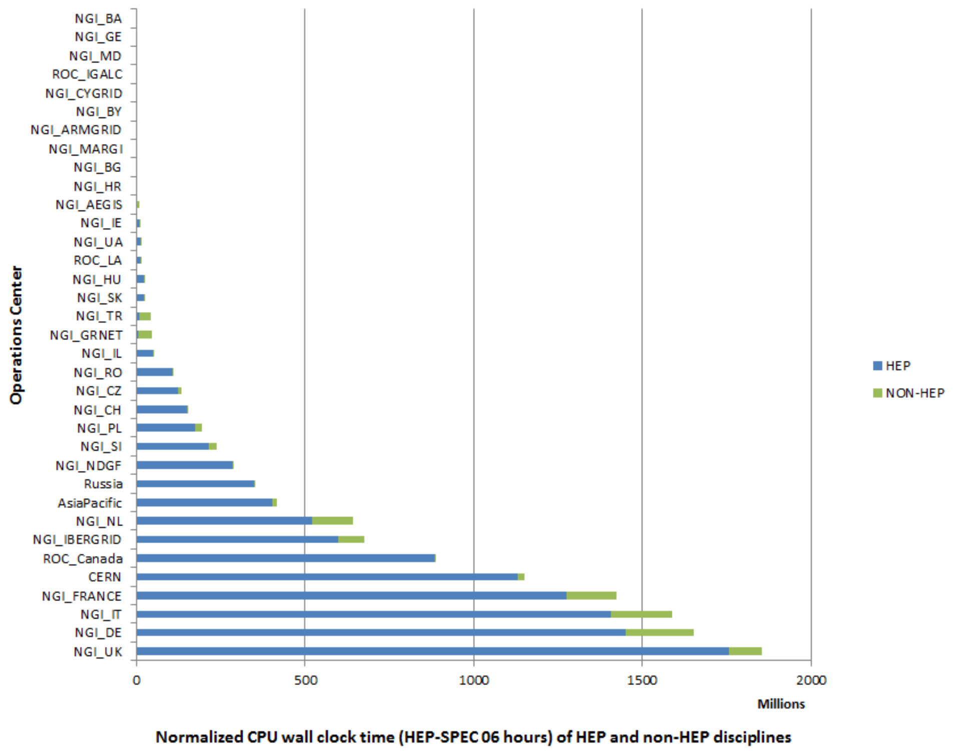

Figure 14. (HEP-SPEC 06 hours) from May 2012 to April 2013 (source: accounting portal). HEP usage is displayed in blue while the aggregated usage of non-HEP disciplines is in green

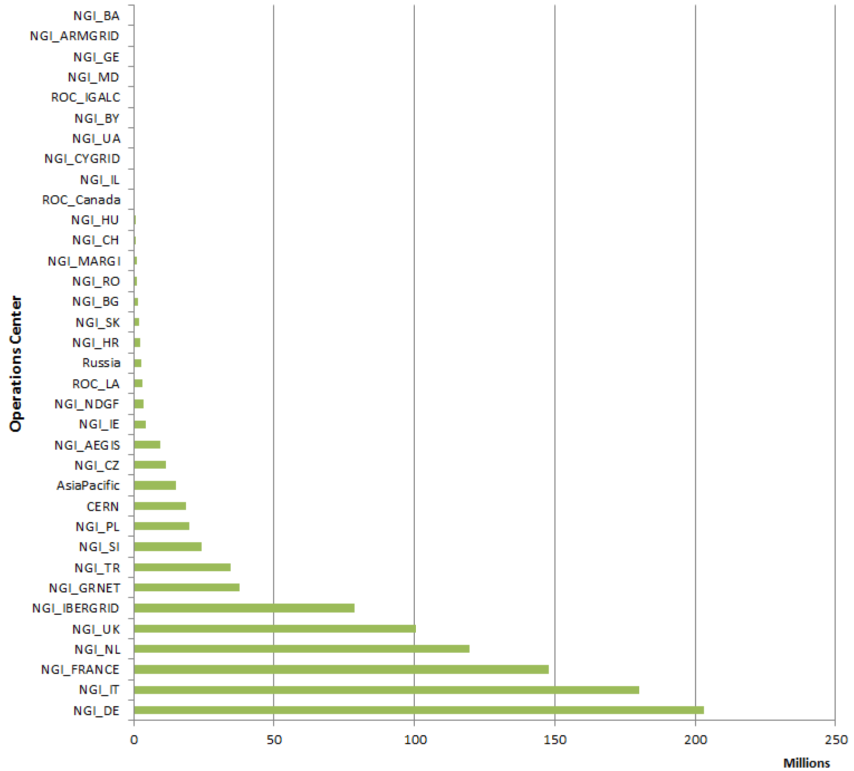

Figure 15. Distribution across EGI Operations Centres of aggregated usage of non-HEP disciplines (CPU wall clock time in HEP-SPEC 06 hours) from May 2012 to May 2013 (source: accounting portal).

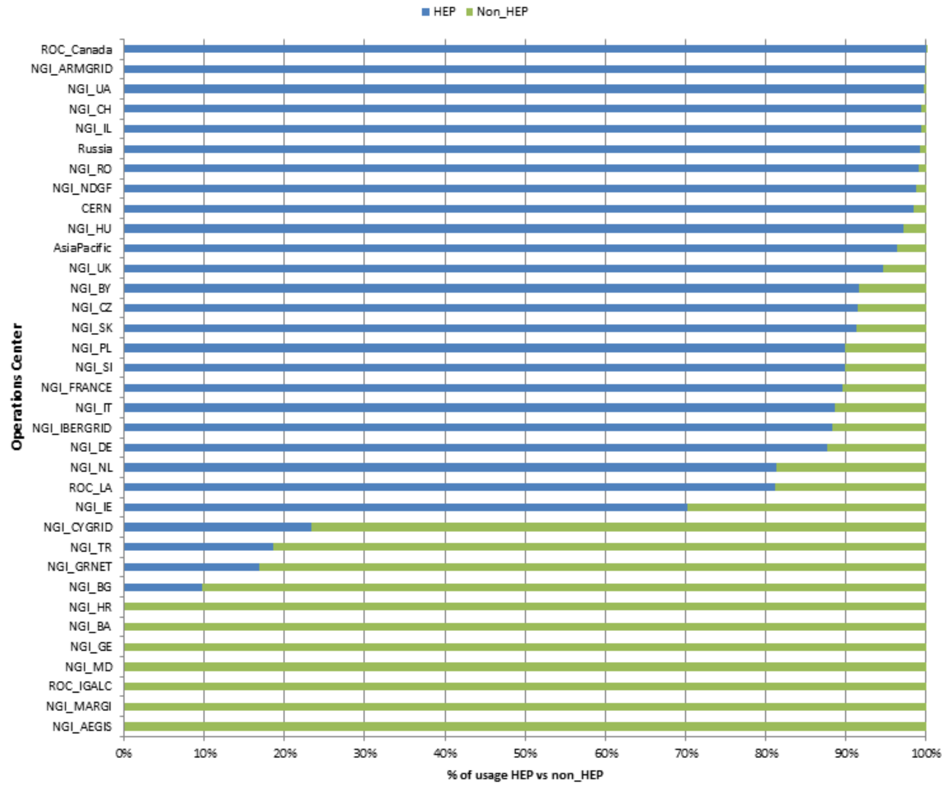

Figure 16. Distribution of resource usage (%) across HEP and non-HEP disciplines from May 2012 to April 2013 (source: accounting portal).

6 SERVICE LEVELS

Services are monitored at three different levels:

- Resource Centre Services;
- Resource infrastructure Provider Services
- EGI.eu Services.

For each category a different set of service level and targets are defined and periodically reviewed. For each set of service levels various reporting systems are available, and are detailed in the following section. The service levels and targets – summarized in 6.1, are formally defined in the RC Operational Level Agreement [RCO], in the RP Operational Level Agreement [RPO] and EGI.eu Operational Level Agreement [EGIO]. EGI has also started work on calculation of VO availability and reliability metrics.

6.1 Service Level Targets and Reporting

This section provides a summary of the EGI operations service level targets formally agreed between resource providers, and periodically reported on a monthly basis.

- Resource Centres[15]
    - Minimum Availability: 70%
    - Minimum Reliability: 75%
    - Reports: https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Resource_Centres
- Resource infrastructure Providers[16]
    - Minimum top-BDII Availability: 99%
    - Minimum top-BDII Reliability: 99%
    - Maximum Regional Operator on Duty Performance Index (see section 6.4.1): 10
    - Reports: https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Resource_infrastructures_Providers
- EGI.eu[17]:
    - Depending on type of service different service targets were defined:

[15] https://wiki.egi.eu/wiki/SLM_RC_Service_Levels

[16] https://wiki.egi.eu/wiki/SLM_RP_Service_Levels

[17] https://wiki.egi.eu/wiki/SLM_EGI.eu_Service_Levels

Table 10 EGI.eu Service Level Targets

Type of service	Service	Service Target
Consulting and Support	1st , 2nd and 3rd Level Support	Top priority: immediate within support hours Very urgent: within 8 support hours Urgent: within 16 support hours Less Urgent: within 40 support hours
	Grid Oversight	1 hour response time within the support hours
	Central EGI helpdesk	Availability/Reliability: 99%/99%
Software	Repository of validated software	Availability/Reliability: Repository frontend: 90% / 99% Repository backend: 90% / 99%
Operations Tools and Services	Service Availability Monitoring (SAM) central service	Availability/Reliability: 95%/99%
	Operational Tools and Meta-service Monitoring (Ops-Monitor)	Availability/Reliability: 99%/99%
	Operations Portal	Availability/Reliability: 99%/99%
	Accounting Portal and database	Availability/Reliability: Repository: 99%/99% Accounting Portal: 99%/99%
	GOCDB	Availability/Reliability: 99%/99%
	Security monitoring tools	Availability/Reliability: 99%/99%
	Grid Services for RC certification	Availability/Reliability: 99%/99%
Grid Services	Virtual Organisation Management	Availability/Reliability: 99%/99%
	Workload Management	Availability/Reliability: 99%/99%
	Information Discovery	Availability/Reliability: 99%/99%

- EGI.eu service level reports are currently under development within JRA1 and are expected to be available from the Operations Portal from June 2013.

6.2 RC Performance

6.2.1 Availability and Reliability

Table 11. EGI-wide Availability and Reliability and the related project metric target.

EGI Average Monthly Reliability	May 2011-January 2012	Y3 Target
Reliability	94.86 %	95%
Availability	93.74 %	-

The quality of grid services deployed by Resource Centres is being measured since 2008 with availability and reliability metrics, computed from the results of periodic tests performed at all certified centres through the Service Availability Monitoring framework (SAM) [SAM]. Availability and reliability metrics were defined to quantitatively express the level of functionality delivered by grid services to end-users with the ultimate goal of identifying areas of the infrastructure needing improvement.

The capability of closely reflecting the experience of the end-user depends on the tests performed. In order to correctly mimic user workflows user-specific tests can be run by customized user-specific

SAM installations [SAMV]. The EGI monthly availability and reliability reports are based on tests (run using the OPS VO), which are sufficiently generic to allow a comparison across all Resource Centres of the infrastructure.

Availability of a service (or a site, depending on the level of aggregation) represents the percentage of time that the services (or sites) were up and running ([uptime / total time] * 100), while Reliability is the percentage of time that the services (or sites) were supposed to be up and running, excluding scheduled downtime for maintenance and other purposes ([uptime / (total time – scheduled down time)] * 100) [AVL].

Certified Resource Centres guarantee 70% availability and 75% reliability for their services. The minimum availability and reliability values accepted for a Resource Centre are defined in Operational Level Agreements established with EGI.eu.

Increasing the overall performance delivered to users has been an on-going effort since the introduction of service level management. Availability/Reliability averaged per quarter across the whole infrastructure have been both steadily increasing from 2008 by approximately 1% per year, moving from 91.9%/93.3% during May 2009 – April 2010 (last year of EGEE-III), to 94.50%/95.42% during May 2011 – April 2012 (second year of EGI-InSPIRE). In the period May 2012 – February 2013 this slightly decrease to 93.74%/94.86%. This is probably related to the maintenance work at RCs to upgrade their software infrastructure (see section 7.2). The PY3 Reliability target (95%) was not met with a small deviation (-0,14%) – see Table 11.

The trend of the overall EGI RC availability and reliability is shown in Figure 17, which plots the average quarterly availability and reliability of RCs from May 2010. For example, the drop recorded from August to November 2011 reflects the transition of various large federated Operations Centres evolving towards a set of operationally independent NGIs.

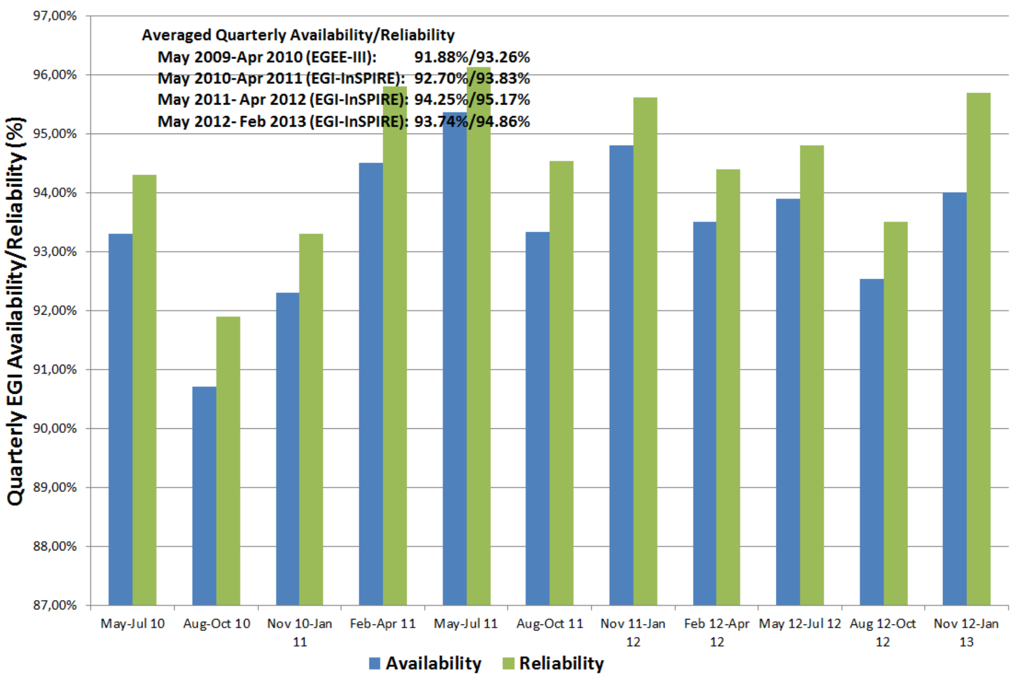

Figure 17. Quarterly availability and reliability of resource centres averaged across EGI from May 2010 to end of PQ7. Source: Availability and reliability monthly reports.

Since November 2012 the responsibility to follow up underperforming RCs was handed off from Grid Oversight team to Regional Operator on Duty teams. This change allows for a quick automated notification to be sent to the RC administrators through the Operations Dashboard in case of a RC failing to meet the minimum requested performance level. A specific Nagios probe was developed for this and as of November 2012 ROD teams are requested to follow the "Quality verification of monthly availability and reliability statistics" procedure[18] for handling of performance issues.

Starting with PY4 RC availability and reliability statistics will be complemented by a new set of VO-oriented availability and reliability statistics, which will more accurately represent the performance perceived by VOs when using the distributed EGI services.

6.3 RP Performance

Table 12. Yearly average availability and reliability of NGI functional services (May 2012-March 2013)

NGI Services Average Monthly Performance	May 2012-March 2013	PY3 Target
Reliability (MSA1.Operations.4)	99.98 %	97%
Availability	98.17 %	-

The performance experienced by users not only depends on resource-access services provided by the RCs, but also on other top-level collective grid services operated by NGIs/EIROs. For this reason, in September 2011 the performance measurement framework was extended to include the core grid services operated by the NGIs and accredited by them to provide access to distributed resources.

RP performance is reported monthly. The purpose of this reporting is to check the availability and reliability of core services operated by NGIs and EIROs, which are typically highly critical as these services provide access to RC services, and are often shared across multiple user communities. In order to enhance their robustness and performance, these services frequently comprise distributed physical instances deployed across multiple RCs. In this case, performance results from the compounded availability of the service physical instances.

The average reliability performed by NGI functional services by far exceeded the PY3 target as shown in Table 12.

6.3.1 Availability and Reliability

Current availability and reliability reports include statistics for the information discovery services (top-BDIIs). The set of monitored core services will be extended to include workload management systems, file catalogues, VO management services etc. Topology information about NGI authoritative service end-points is provided by GOCDB through NGI service groups, whose implementation was completed in PY3[19].

It was decided to introduce two profiles for RP availability/reliability calculations:

- NGI OPS profile – monitoring services (SAM, VO SAM) and the regional APEL DB
- NGI Tech profile – other core services

[18] https://wiki.egi.eu/wiki/PROC04_Quality_verification_of_monthly_availability_and_reliability_statistics#Process_of_handling_RC_Availability_and_Reliability

[19] https://wiki.egi.eu/wiki/NGI_services_in_GOCDB

As shown in Figure 18, the performance of NGI services has been excellently improving since January 2012 when the NGI Availability/Reliability statistics were introduced for the first time. As of January 2012, NGIs whose service availability does not reach 99%, are being assisted to define a plan for service improvement. The short term objective of this action, which was the improvement of the performance offered to end-users by NGIs, was successfully accomplished.

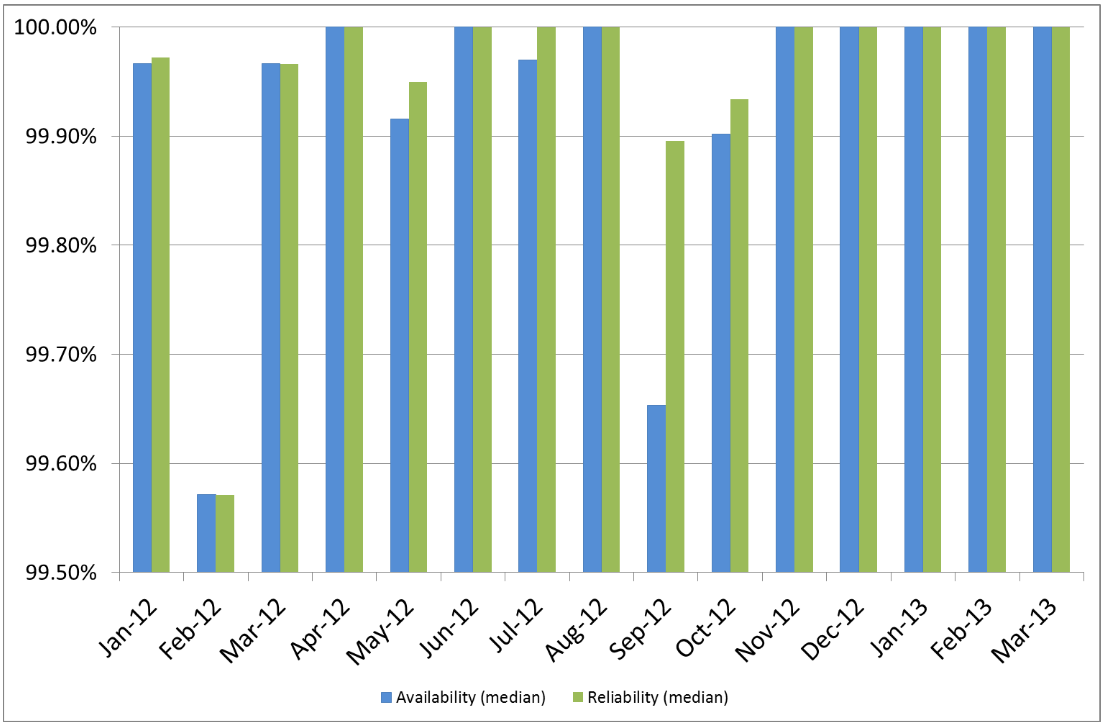

Figure 18. Median of NGI monthly Availability and Reliability performance (top-BDII service) – Jan 2012 – March 2013.

In order to consolidate the information discovery service various actions were undertaken:

- In collaboration with the Distributed Middleware Support Unit, various techniques for the configuration of top-BDII in failover mode were documented in a manual [MAN05].
- The list of authoritative top-BDIIs was collected and their configuration was assessed.
- The list of RCs making use of the CERN top-BDII as primary instance was collected and the NGIs were requested to support the administrators to change configurations, so that the correct authoritative instance is used instead.
- Small NGIs which failed to provide reliable top-BDII can now use the EGI Catch All top-BDII server provided by Greek JRU[20].

6.4 EGI.eu Performance

Table 13. Yearly average availability and reliability of EGI.eu Core Infrastructure Platform (PQ11)

[20] https://wiki.egi.eu/wiki/Catch_All_Grid_Core_Services

EGI.eu Core Infrastructure Monthly Performance (MSA1.Operations.6a)	May 2012-March 2013	PY3 Target
Reliability	98.60%	97%
Availability	98.60%	-

Monitoring of EGI.eu Core Infrastructure Platform was rolled to production in November 2012. The central EGI.eu services being monitored – which are part of the EGI Core Infrastructure Platform – are: the distributed monitoring infrastructure – SAM, the EGI-InSPIRE Metrics Portal, the Accounting Portal and central database, the central Operations Portal and the service registry GOCDB. In order to do so, a new central SAM instance was rolled to production to monitor these tools and various user community services (Training Marketplace, CRM and Application Database).

Availability statistics of these tools are now accessible through the MyEGI portal[21].

The average availability and reliability performed by these EGI.eu tools in PQ11 is indicated in Table 13 and exceeded the PY3 target (97%). The performance was affected by some instability experienced by the accounting portal, the central accounting database and CRM. The performance of these tools is being monitored on a monthly basis. From PQ4 an automated monthly report generator will be available in the Operations Portal. Partners failing to meet the minimum performance requirements defined in the EGI.eu Operational Level Agreement will be requested to provide performance improvement plans.

6.4.1 ROD Performance Index

A performance metric was defined in PY2 to measure the quality of the NGI support services provided by the operations centres. The Regional Operator or Duty team of each operations centre is responsible of monitoring alarms and of proactively contacting site administrators so that the incident is promptly managed (an alarm is generated in case of failure of an OPERATIONS monitoring test).

The ROD performance index[22] is the sum of the number of ticket expired in the operations dashboard daily, and the number of alarms older than 72h appearing in the operations dashboard daily.

The ROD performance index is calculated monthly from the data gathered by EGI Operations Portal, and it does not take into account weekends. The threshold is set to 10 items. Above this value ROD teams have to provide explanations and a plan of improvement of their oversight service.

[21] https://grid-monitoring.egi.eu/myegi

[22] https://wiki.egi.eu/wiki/ROD_performance_index#Definition

Figure 19 shows the monthly number of unhandled items on operations dashboard. The chart shows that the metric is highly affected by holiday periods when ROD teams may not be on duty

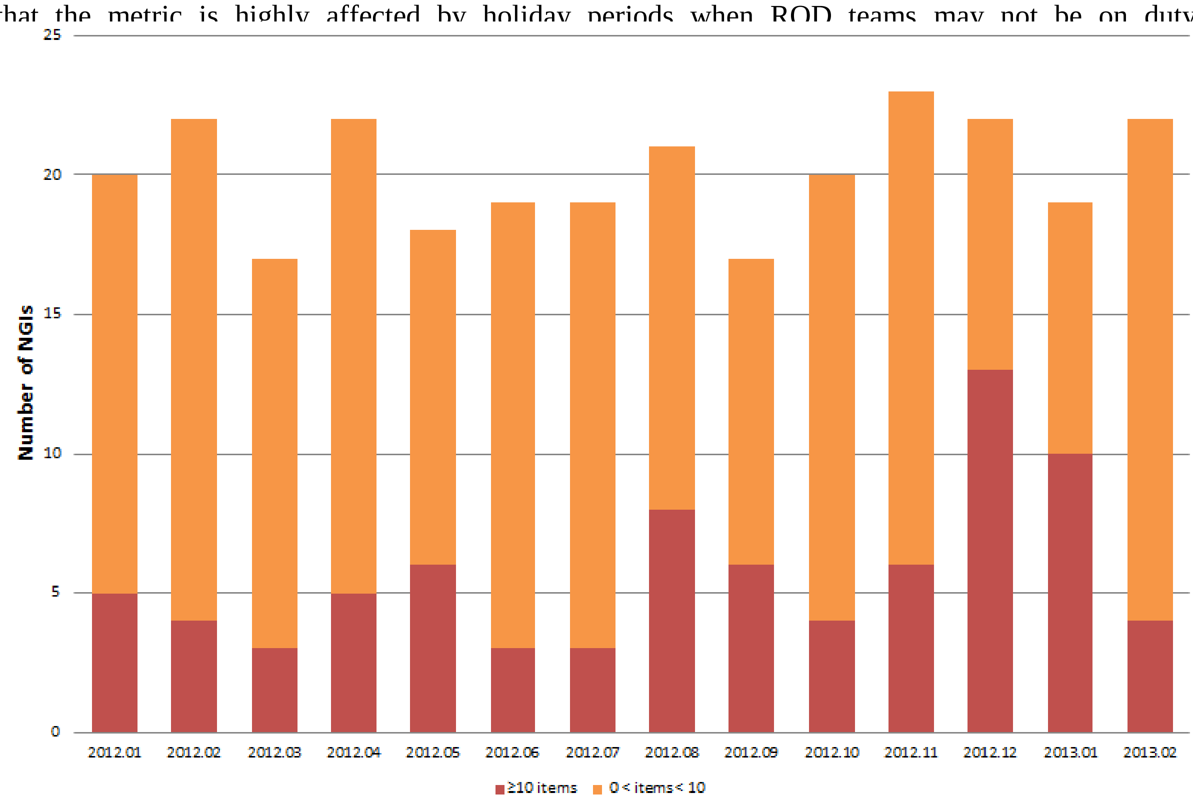

Figure 19. The monthly number of unhandled items on operations dashboard (from January 2012 to PQ11). Source: Operations Portal.

6.5 VO Performance

A new set of Availability and Reliability reports is being developed to provide VO-oriented Availability and Reliability views that only include the service instances on which a given VO is enabled. The list of VO-enabled services is extracted from the information discovery service (top-BDII). Only those services for which this information is published are included in the computation. In the current prototype these are: Compute Element, Storage Element, Local File Catalogue, Workload Management System and VO Membership Service. Results for computation are extracted from monitoring tests run with the OPS VO.

The Availability computation algorithm is such that if the fraction of service instances of a given type scoring 100% on an hourly basis exceeds a given threshold (80% for the results reported in in Figure 20), then the availability of that service group is set to 1, 0 otherwise. This computation is applied on an hourly basis and the aggregation is calculated daily and monthly by averaging the hourly availabilities for each service type.

Figure 20 shows the monthly Availability for April 2013 of the 10 most active VOs (alice, atlas, auger, biomed, cms, compchem, icecube, ilc, lhcb and theophys). The median of the monthly Availability for all VOs is 99.48%, while the average is 97.67%.

This prototype will be validated for all months of PY3 and will be available in production in PQ13 after a testing phase as a module of the Operations Portal.

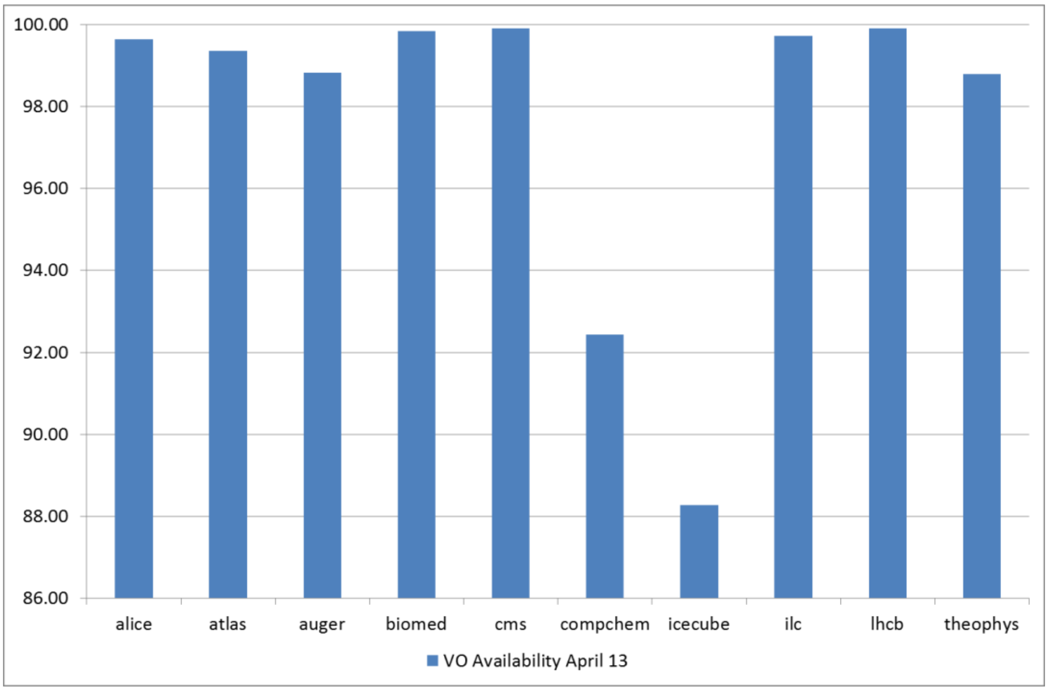

Figure 20. Availability of the top-10 active VOs (April 2013)

7 GRID SERVICES

In this section we review the status of deployment of different software platforms across EGI. As indicated in Table 14, the set of software platforms that are successfully integrated, currently encompasses EMI software (ARC, dCache, ex-gLite products, UNICORE), GLOBUS being maintained, released and supported by the IGE project, QoSCosGrid supported by PL-Grid[23], and Desktop Grid software released and supported by the EDGI project[24]. In PY3 the integration level of the various stacks was consolidated, even though it cannot be considered totally complete yet as accounting integration is still in progress for various platforms. Currently the EGI service registry (GOCDB) defines the service types necessary to register services from all the stacks.

The list of production end-point services per platform can be obtained programmatically from the GOCDB programmatic interface[25].

7.1 Integrated Software Platforms

Table 14. Deployment of integrated software platforms across EGI

Integrated Grid Platform	Number of countries	Countries
ARC	11	Switzerland, Norway, Denmark, Sweden, Finland, Slovenia, Latvia, Germany, Ukraine, Estonia, Lithuania
Desktop Grid (experimental phase)	1	Hungary
GLOBUS	5	GridFtp: United Kingdom, Croatia, Finland, Germany GRAM: Germany, The Netherlands, Croatia,
QosCosGrid (QCG)	1	Poland
UNICORE	2	Germany, Poland

Accounting integration is still in progress for UNICORE, Globus and QCG, while computing resources accessible through ARC-CE and CREAM interfaces have been accounted for their usage from the beginning of EGI-InSPIRE. The Accounting Task Force of the TCB[26] is responsible of leading the extension of the current EGI accounting infrastructure to encompass peer grids and new integrated infrastructures.

Integration of information publishing will be accomplished for ARC and UNICORE in PY4, thanks to the support of GLUE2 information publishing into top-BDII, which is a function that was released for

[23] http://www.egi.eu/community/collaborations/MAPPER.html

[24] http://www.egi.eu/community/collaborations/EDGI.html

[25] https://wiki.egi.eu/wiki/GOCDB/services

[26] https://wiki.egi.eu/wiki/TCB:Accounting_Task_Force

deployment by EMI with release v3. Information publishing for the GLOBUS GRAM service was made available by IGE with release v3.2[27].

The services originating from the gLite distribution (now unsupported) and now supported and distributed through EMI releases, are deployed by the majority of the production RCs. However, the number of operations centres supporting other stacks slightly increased during PY3: as shown in Figure 21, seven operations centres are deploying ARC middleware, namely: NGI_NDGF (including Denmark, Estonia, part of the Finnish resources, Latvia, Norway, Sweden, Lithuania), NGI_CH, NGI_DE, NGI_FI, NGI_SI and NGI_UK.

UNICORE is supported by two operations centres: NGI_DE and NGI_PL, while Globus middleware is deployed by NGI_DE, NGI_FI, NGI_HR, NGI_NL and NGI_UK. QosCosGrid middleware is deployed only by NGI_PL.

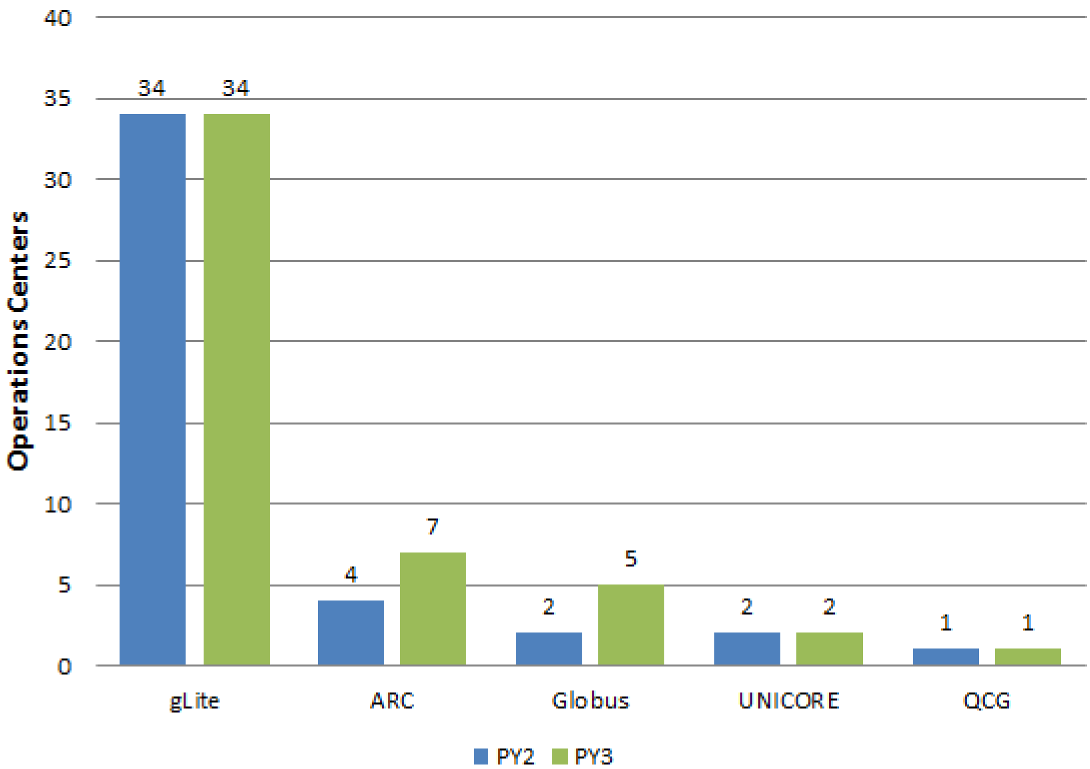

Figure 21. Deployment of the five reference grid middleware stacks across the EGI-InSPIRE operations centres, March 2013 (source GOCDB).

[27] https://sites.google.com/a/ige-project.eu/public/downloads/software/releases/320

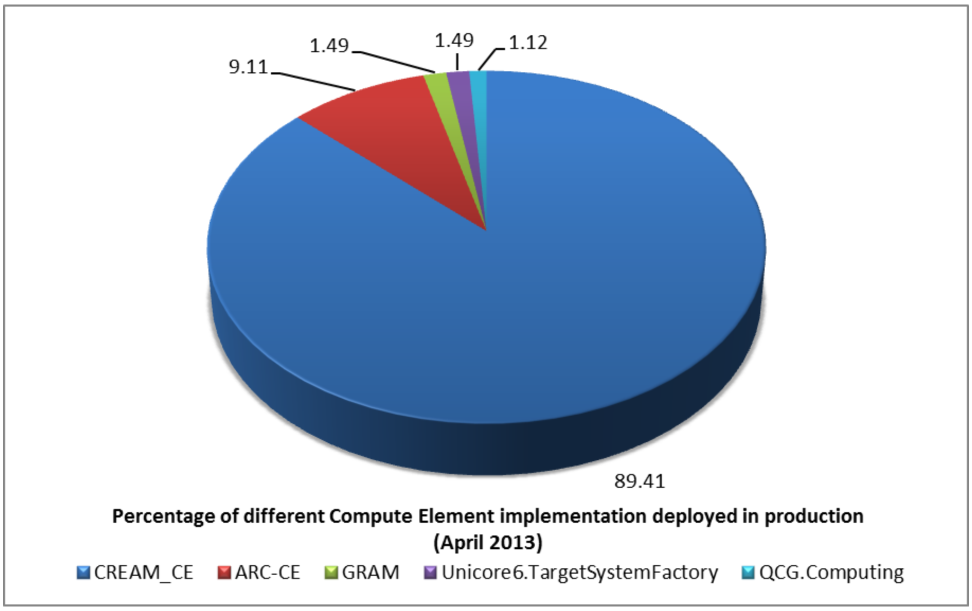

Figure 22. Number of instances of the different implementations of the compute capability, across the EGI_InSPIRE partners and the integrated Resource infrastructure Providers, March 2013 (source: GOCDB)

Various middleware stacks are in production in EGI. An indication of their distribution is given by the various Compute Element deployed by Resource Centres. Figure 22 shows this distribution: CREAM-CE is in production in the 89.41% of the infrastructure, ARC-CE is second in deployment (9.11%) followed by GRAM (1.49%), Unicore6.TargetSystemFactory (1.49%) and QCG.Computing (1.12%).

LCG-CE reached end of support at the end of April 2012. All LCG-CE instances were successfully upgraded (the majority by the end of 2012) to other supported CE implementations.

7.2 Software retirement

In PQ10 a new policy for the retirement of unsupported software from the production infrastructure was approved. This policy was incorporated into the main body of EGI security procedures and new procedures were developed to support the timely retirement of software[28]. The policy says that in compliance to the EGI Service Operations Security Policy [SOSP][29], unsupported software SHOULD be decommissioned before its End of Security Updates and Support, and MUST be retired no later than one month after its End of Security Updates and Support. After this date, if a critical vulnerability were to emerge in the software, EGI CSIRT can request the service to be turned off immediately.

The main software distributions being deployed to date in EGI are EMI[30] and IGE[31].

---

[28] https://wiki.egi.eu/wiki/PROC16 and https://wiki.egi.eu/wiki/PROC01.

[29] A Resource Centre Administrator SHOULD follow IT security best practices that include pro-actively applying software patches, updates or configuration changes related to security.

[30] http://www.eu-emi.eu/releases

[31] https://sites.google.com/a/ige-project.eu/public/downloads/software/releases

gLite 3.1 and 3.2 distributions are no longer supported. The decommissioning campaign of these two distributions started in October 2012. This first decommissioning campaign was subsequently followed by an EMI-1 decommissioning campaign which is still in progress to date.

Software decommissioning involved EGI.eu operations, EGI CSIRT, the Security Policy Group (for the definition of a software retirement policy) and the Central Grid Oversight time for the enforcement of retirement policies across the whole infrastructure. In addition, to streamline software retirement and monitor progress, the security monitoring team developed and deployed new custom security probes as required for monitoring for deployed software beyond end of support. A dedicated Nagios service[32] was deployed to monitor middleware components.

gLite 3.2 was supported until the end of April 2012 and many sites migrated from gLite3.1 directly to the functionally equivalent EMI-1 components (supported until April 2013) or to the upcoming EMI-2 release supporting both SL5 and SL6.

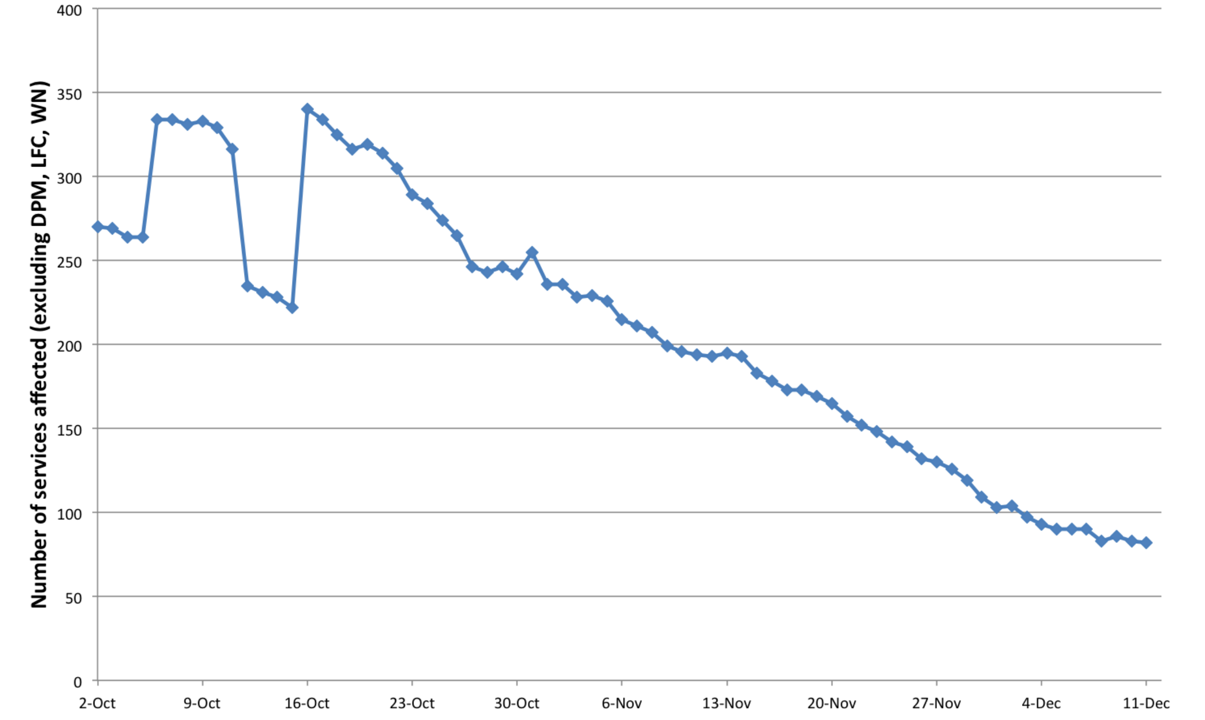

Figure 23. Number of unsupported gLite software services deployed in production. This diagrams shows the progress of the gLite decommissioning campaign in PQ09, 10 and 11.

Figure 24 shows the progress of decommissioning gLite 3.1 and 3.2 unsupported middleware, which was successfully completed in PQ11. The increase in the number of service instances to be decommissioning, which is visible in October 2012, is due to the introduction of new Nagios probes for the automated detection of new unsupported software versions.

7.3 Core Middleware Services

Core grid middleware services are provided by Resource infrastructure Providers to fulfil the needs of the national and international VOs supported by their resource centres. There are many core services

[32] https://midmon.egi.eu/nagios/

provided through the different middleware stacks, this paragraph provides a snapshot of the current deployment for the four most deployed ones: LFC (file catalogue), WMS (workload management), Top-BDII (information system top-level cache) and VOMS (VO membership, attribute management).

Figure 25 shows the current distribution of production instances among the EGI-InSPIRE partners and integrated resource providers. The instances information was collected by querying the Top-BDII: this information source contains also the software version which is not available in the services registry (GOCDB).

As of March 2013 the EGI integrated infrastructure comprises 367 core services: 66 VOMS instances, 160 WMS, 33 LFC and 108 Top-BDII. The number of core services operated by an NGI naturally grows with the number of sites, the number of user communities supported and the size of the supported VOs.

WMS is the service with the highest number of instances, often NGIs deploy multiple instances of WMS to load balance the workload on individual service instances.

The Top-BDII is offered – either directly or through the provisioning by other partners – by all the NGIs who are deploying EMI middleware, since it offers a critical capability for service discovery.

<table>
<tr>
<td>(a) Overall distribution</td>
<td>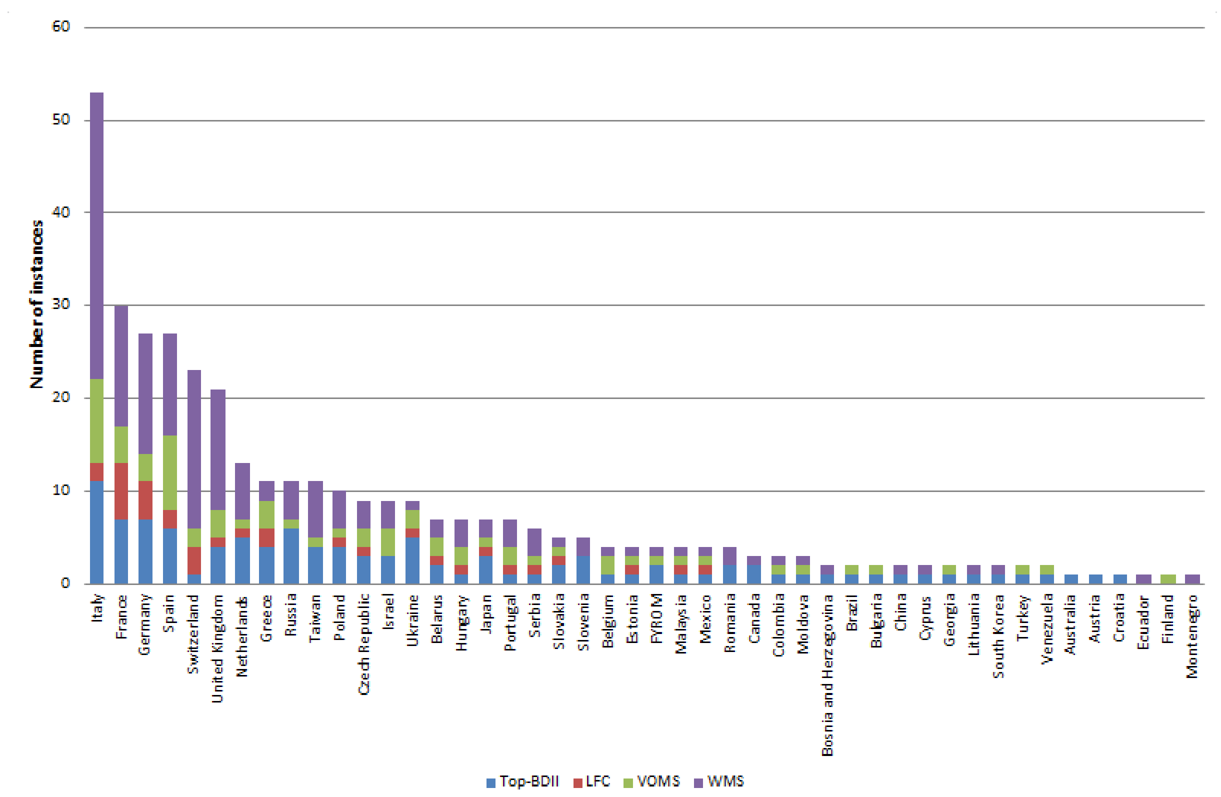
</td>
</tr>
</table>

(b) Workload Management (WMS)	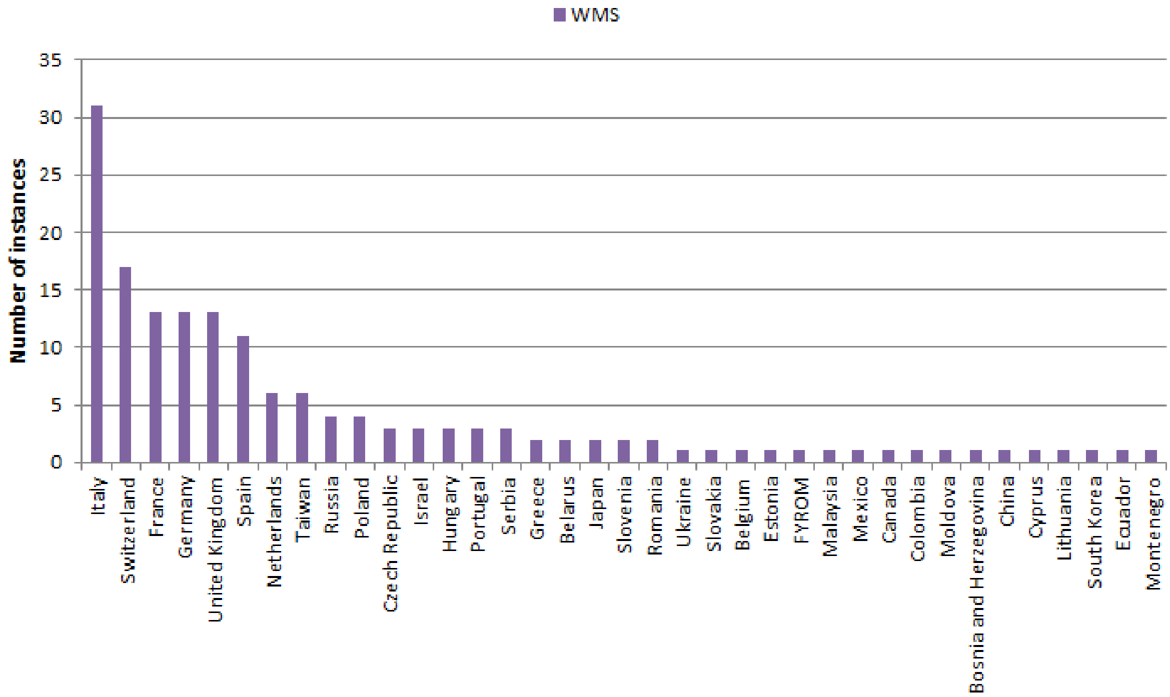
€ VO membership (VOMS)	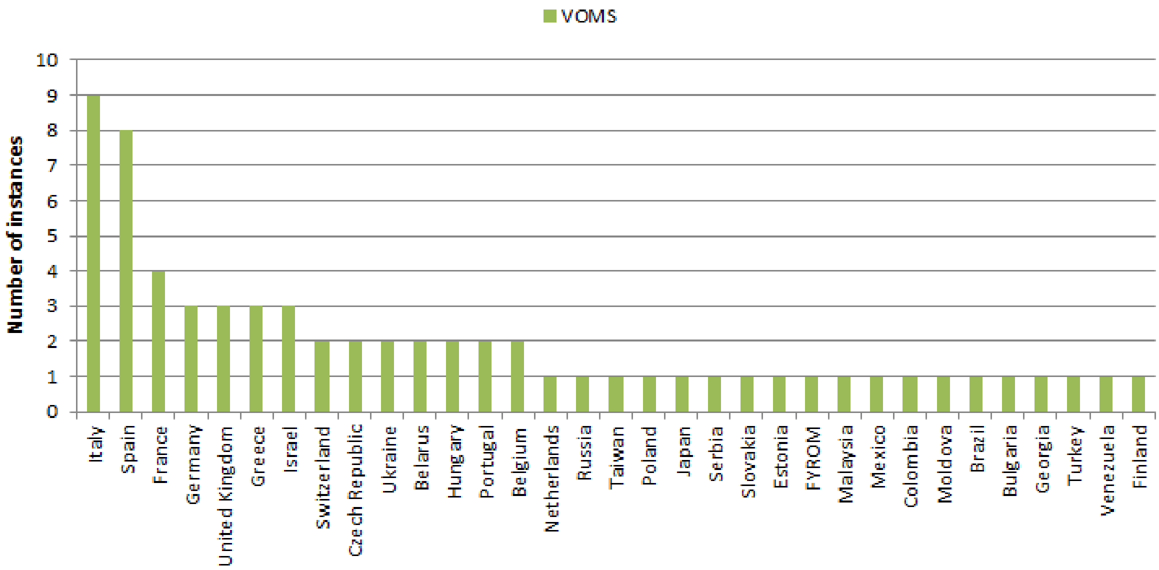

(d) file catalogue (LFC)	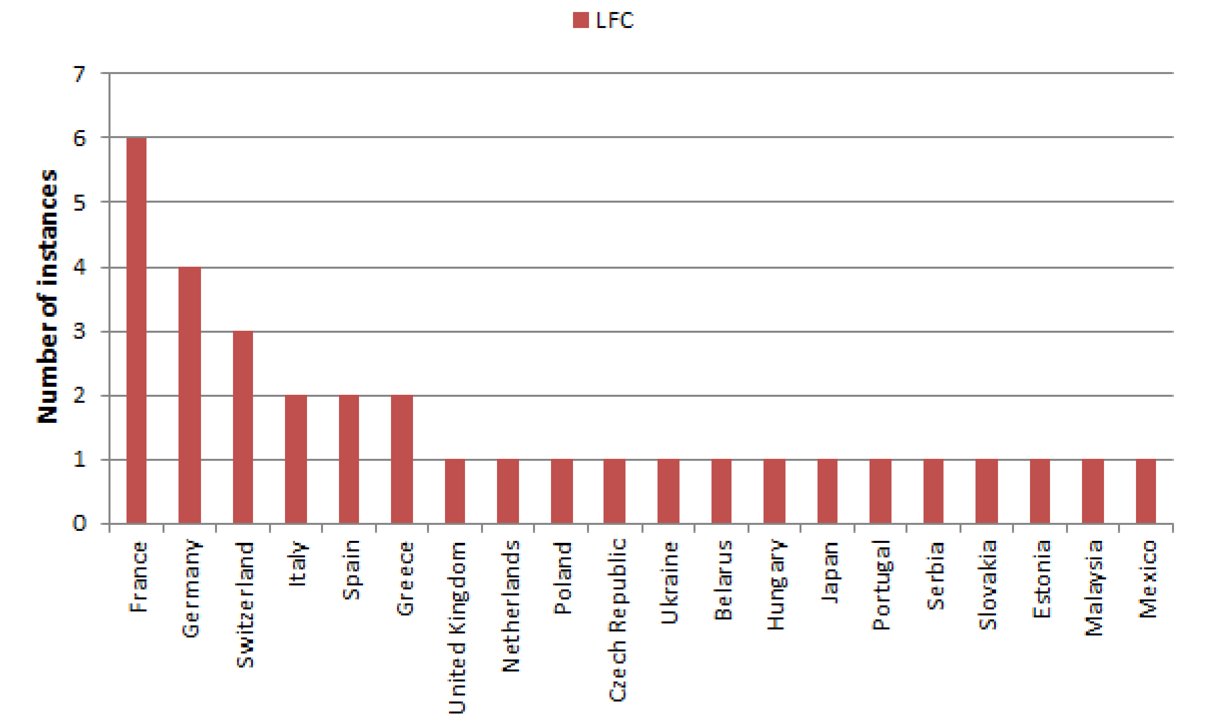
€ Information system (Top-BDII)	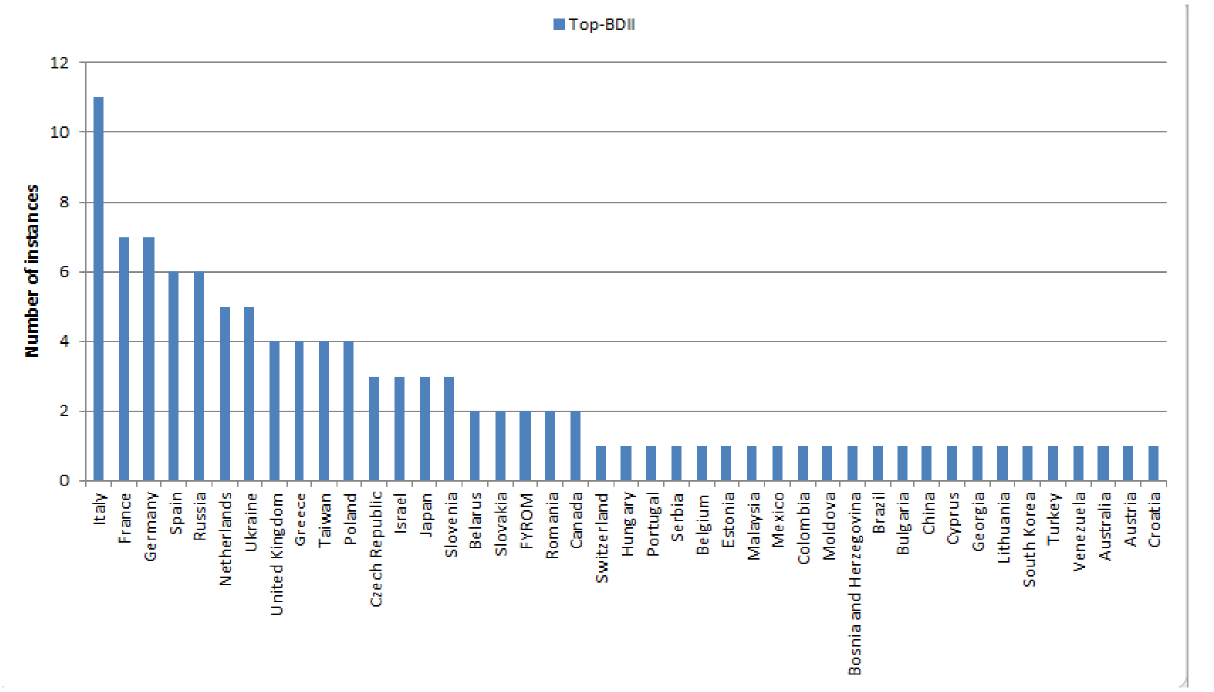

Figure 25, Number of core services instances deployed within the EGI-InSPIRE integrated infrastructure. Source: GOC DB (March 2013).

7.4 UMD in PY4

The UMD distribution will continue in PY4 to be the recommended distribution for RCs including products that successfully passed verification and validation according well defined procedures[33].

With the end of EMI and IGE and their coordination function in May 2013, a number of actions have been undertaken to adapt EGI processes and support structures to the changes introduced.

Support. EMI and IGE are the projects that have been providing specialized 3rd level support to EGI users and operators and internal coordination of 3rd level support. With the end of EMI and IGE Product Teams will be responsible of supporting their own products according to an own set of software support policies. In PY12 all Product Teams where contacted to discuss future support channels, response time to incident records in GGUS and Support Unit structures.

In order to simplify software support, a GGUS workflow was defined and approved to handle unresponsive Support Units, with the objective of making sure that supporters are periodically notified when a response is due[34].

The current technology helpdesk for 3rd level support was assessed and a number of changes were agreed to improve accessibility to information in tickets that reach the 3rd level support escalation stage, and to improve traceability of delivery dates of Requests for Changes.

The new Support Unit structure, the workflow for unresponsive supporters and the workflow for handling 3rd level support tickets will be implemented in PY13. 1st and 2nd level support, which are services internally provided by EGI, are not affected by changes introduced in 3rd level support.

Coordination of UMD releases. To compensate for the discontinuation of several coordination functions currently supplied by EGI and IGE as of May 2013, a new board for coordination of activities for UMD release activities was defined. The board – called UMD Release Team (URT) – will provide lightweight coordination of Product Team release activities. The board will discuss issues found during verification and validation of software, release calendars for fixes to critical issues affecting the infrastructure, UMD quality criteria, and will provide information about UMD release calendars and a communication channel between Product Teams. Product Team representatives and the UMD software provisioning team are members of URT. The URT was constituted in PQ12 and the Terms of Reference are being defined [URT].

Future support of software products. With the end of EMI the continuation of software maintenance and development activities for some products may be compromised. Continuity of support of several products is being discussed with the relevant Product Teams[35].

---

[33] https://wiki.egi.eu/wiki/EGI_Software_Provisioning

[34] https://wiki.egi.eu/wiki/FAQ_GGUS-Waiting-For-PT-Process

[35] https://twiki.cern.ch/twiki/bin/view/EMI/EmiProductTeams

8 CLOUD INFRASTRUCTURE PLATFORM

To provide generic, consistent and flexible access to EGI resources, EGI initiated a strategic activity to establish a federation of locally deployed IaaS Clouds. The EGI Cloud Infrastructure Platform directly supports EGI's strategic alignment with the European Commission's Horizon 2020 strategy. While EGI will continue to support and maintain its existing relationships with research communities, the Cloud platform will be offered in support of new research communities stemming from the so-called "long tail of science". In compliance with the Cloud computing model, the EGI does not mandate deploying any particular or specific Cloud Management stack; it is the responsibility of the Resource Providers to research, identify and deploy the solution that fits best their individual needs for as long as the offered services implement the required interfaces and domain languages.

Consequently, the EGI Cloud Infrastructure Platform is built around the concept of an *abstract* Cloud Management stack subsystem that is integrated with components of the EGI Core Infrastructure Platform (CLIP), that are necessary to federate Distributed Computing Infrastructures into a (set of) consistent resource access services across administrative domains (nationally or globally). The different cloud management middleware are federated by providing common interfaces to access the virtualized resources, such as Open Cloud Computing Interface (OCCI) and Cloud Data Management Interface (CDMI).

Table 15: Resource providers participating to the Federated Clouds test-bed (April 2013)

Resource centre name	Number of cores available in the test-bed	Amount of disk space available in the test-bed	Cloud middleware deployed
BSC	96	3.6TB	Open Nebula/Open Stack
CESGA	33	450GB	Open Nebula
CESNET	240	44TB	Open Nebula
Cyfronet	32	Local disk	Open Nebula
FZ Jülich	76	5TB	Open Stack
GRNET	200	22TB	Okeanos (Open Stack compatible)
GWDG	32	1TB	Open Nebula
CC-IN2P3	384	32	Open Stack
KTH	4	1TB	Open Nebula

Table 16: Resource providers under the process of integration in the Federated Clouds test-bed (April 2013)

	Number of cores potentially contributed to the testbed	Amount of disk space potentially contributed to the test-bed	Cloud middleware deployed
CETA-CIEMAT	104	Local disk	Open Stack
INFN	24	2TB	WNoDeS
IFCA	256	Local disk	Open Stack

Table 15 contains the resource providers participating to the activities of the Federated Cloud task force. As shown in the list the most common cloud middleware solutions are Open Nebula and Open Stack. In addition, some NGIs are deploying cloud management software developed within their organization such us Okeanos and WNoDeS. Ten Resource Providers are being monitored with a test instance of SAM: the monitoring service is an instance of the SAM production distribution, with in addition a set of cloud-specific probes. All the resource providers monitored are also registered in GOCDB. Table 16 contains the list of resource providers under the process of integration; the representatives of these resource centres are already participating to the task force activities, and their resources will be part of the test-bed in the coming months. The two tables reports a summary of the resources contributed to the test-bed by the resource centres in terms of number of physical cores and disk space, for some resource providers the disk is only available as local disk in the machines used to run the virtual machine instances.

The EGI Federated Cloud testbed will be integrated into the production infrastructure in PY4.

9 STAGED ROLLOUT INFRASTRUCTURE

In a large-scale distributed infrastructure, deployment of software updates requires coordination and needs to follow a well-defined process. In EGI this is implemented by gradually installing updates that successfully passed internal verification, in a selected list of Resource Centres. This process is called *Staged Rollout* and the Resource Centres performing the function of tester, are named *Early Adopters* (EAs) [SRW]. The Staged Rollout services hosted by the EA Resource Centres constitute together the *Staged Rollout Infrastructure*, which is distributed as Staged Rollout and is a joint effort of the EGI Operations Community.

The process aims at collecting information about the performance of a new software release when deployed in a production environment: this includes checking installation and configuration, as well as functionality, robustness and scalability of the software especially when interworking with other Grid services as required in real user workflows. The successful Staged Rollout of software is a precondition for declaring it ready for deployment. This process is coordinated by EGI.eu to ensure a successful and tight collaboration between the various stakeholders: Resource Centres, Technology Providers, the EGI.eu technical management and the EGI repository managers.

EAs are not testers responsible of software certification, as software distributed through the Unified Middleware Distribution [UMD] is certified by the Technology Providers. Software under validation is accessible from a specific dedicated software repository. For information about the UMD Software Provisioning Process see [MS512].

Table 17. Overview of EGI-InSPIRE Staged Rollout metrics.

Metric	PQ8	PQ9	PQ10	PQ11
Number of staged rollouts	12	78	40	48
Number of components	8	54	29	32
Number of sites	9	24	22	20

As shown in Table 17 the largest number of products were tested in PQ9 in preparation to the release of the Unified Middleware Distribution 2. This number was gradually reduced in the following quarters following the release schedule of EMI and IGE updates. The number of participating EAs has been progressively increasing to test a growing set of products from EMI, IGE and EGI-InSPIRE JRA1 (operational tools), and it currently amount to 74 teams.

The staged rollout of QosCosGrid software is expected in PY4 in preparation to the Unified Middleware Distribution release 3.

The number of tests performed from PQ8 to PQ11 by NGIs and EIROs is plotted in Figure 26.

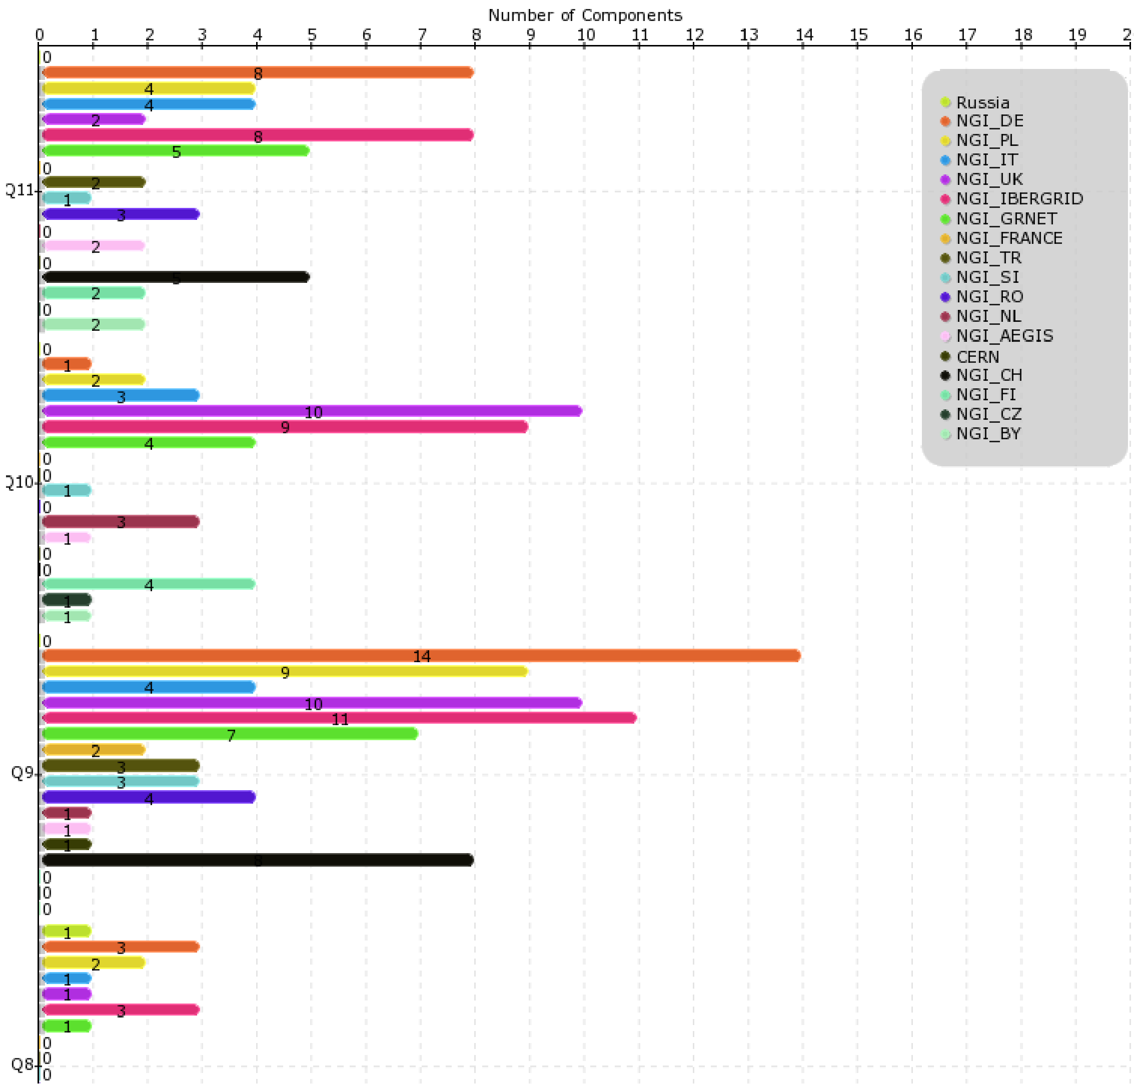

Figure 26. Number of Staged Rollout tests performed from PQ8 to PQ11 by NGIs/EIROs. (source: Staged Rollout portal)

10 SOFTWARE SUPPORT

Software support in PY3 followed the procedures of Deployed Middleware Support Unit (DMSU) established in PY2. However, at the end of PY2 changes were proposed to merge the former TPM activity (in TSA1.7) and DMSU (TSA2.5) into a single task in SA1, with the main goals to avoid duplicating work on receiving and assessing software tickets, and to optimize the task workflows. The analysis was done in the "Revision of TPM and DMSU activities" [36] document, its outcome – the desired state was described in "MS511 Deployed Middleware Support Unit Operations Procedures" [37], and after having been approved by the project review, the changes were implemented in early autumn 2012. The following main activities for software support process were identified:

- Ticket triage and assignment for dispatching of tickets to the appropriate SUs within GGUS
- 2nd level software support, encompassing both grid middleware and operational tools
- Ticket oversight and follow-up

According to this split of responsibilities, the roles were reassigned to the involved partners of the former TPM and DMSU, and the project effort assignment was slightly adjusted. In particular, the coverage was extended to support EGI operational tools and other products. The ticket payload of the 2nd level support unit followed the trends of the former DMSU.

The following table shows the number of software support tickets handled in PY3. The number of these (619 tickets in total) is lower with respects to the same period of PY2 (730). This can be explained by the rather high number of tickets related to the pre-release testing of UMD 1.0.0.

Table 18. Number of software support tickets handled in GGUS (1^{st,} 2nd and 3rd level support)

Metric	PQ9	PQ10	PQ11	Feb-Mar 2012
Assigned to DMSU	179	156	173	111
Reassigned to TPM	22	8	6	2
Reassigned to 3rd level	116	105	130	76
Solved by DMSU	40	48	52	32
Mean/median days to solve	18.8/10.8	28.5/11.1	19.1/4.0	22.5/4.7

The overall ratio of tickets solved by the support unit is 27%, which is a clear improvement compared to the rate accomplished in PY2 (21%). The process of analyzing tickets before reassigning them to 3rd line support units (deployed in PY2 and followed throughout PY3), contributed to this improvement.

The following graph shows the weekly distribution of tickets. Oscillations in load are considered to be normal. The workload reduces – as expected – in the summer and Christmas period.

[36] Revision of TPM and DMSU activities, https://documents.egi.eu/document/1104

[37] MS511 Deployed Middleware Support Unit Operations Procedures, https://documents.egi.eu/document/1134

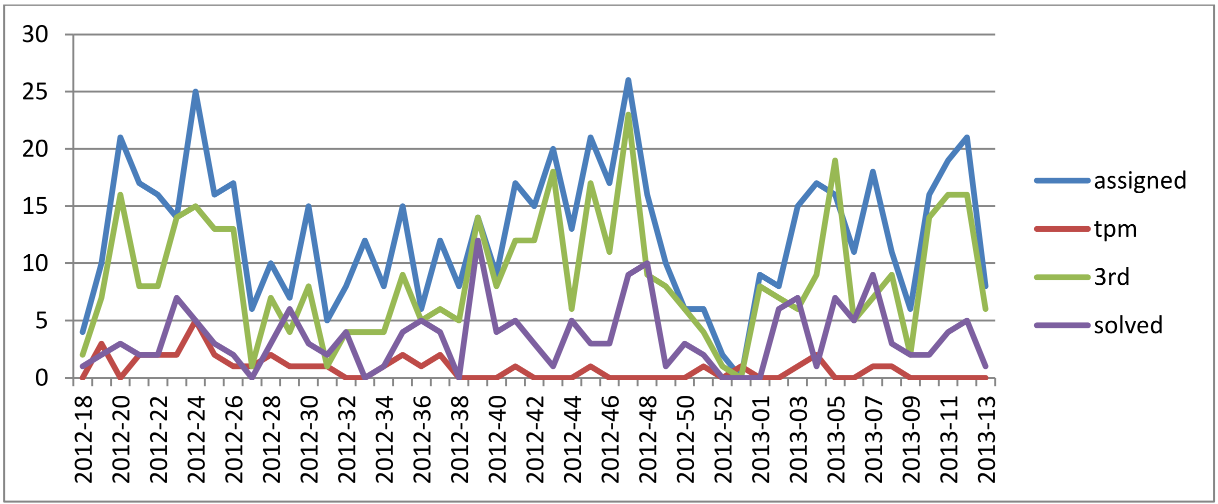

Figure 27. Weekly distribution of the software tickets handled by the 2nd level Support Unit.

Out of those total number of tickets, only 2 were *top-priority*, while 25 were assessed to be *very urgent* (the two highest priority levels according to the GGUS classification). These are reasonable numbers in which the special treatment of the tickets – requiring negotiation with the 3rd line support team to ensure that those tickets are handled in a timely manner – can be considered to be feasible.

The software support unit interacts with the EGI Operations on a regular basis. Issues that are identified to have a potential broader impact on the infrastructure, are described in a dedicated wiki page[38]. The unit representative also attends the regular bi-monthly operations meetings where those issues are discussed, and eventually further issues are fed back to the software support unit. The unit leader also attends the TCB meetings.

Starting in 2013 further adjustments to the ticket follow-up process are being discussed and will be implemented in PQ13. These are needed to adapt to the discontinuation of the software support coordination function currently provided by the EMI and IGE EC project.

[38] https://wiki.egi.eu/wiki/DMSU_topics_gridops_meeting

11 CONCLUSIONS

The production Infrastructure satisfactorily met the PY3 targets of the SA1 project metrics: the number of RCs integrated, number of job slots offered, and the usage. The Desktop Grid integration is being piloted in Hungary. The contribution of EGI-InSPIRE SA1 to the accomplishment of the applicable project objectives is described[39].

- Objective 1 (O1): *The continued operation and expansion of today's production Infrastructure*.

  This objective was successfully met by completing the integration of the Ukrainian National Grid[40] comprising 12 production RCs. A MoU with the *Asia Pacific Grid Initiative* (APGI) was signed in PQ12 and a MoU is being finalized with Open Science Grid in USA. Unfortunately two Operations Centres were decommissioned because of sustainability issues: NGI Ireland and Iniciativa de Grid de America Latina – Caribe. Fortunately this was compensated by a substantial increase in the offered capacity: compute resources increased by +33.6% in PY3, while disk capacity increase to 177 PB (+25.36%). At the end of PQ11 the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to 347,307, which provide 3.32 Million HEP-SPEC 06. The performance of NGI services has been excellently improving since January 2012 when the NGI Availability/Reliability statistics were introduced for the first time, and the EGI Core Infrastructure Platform is delivering very good and stable performance.

- *Objective 2 (O2): The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.*

  In PY2 the responsibility of providing VO services was migrated to the EGI.eu operations team and the NGIs. VO support includes existing SA1 VO services provided by NGIs including support through the EGI helpdesk, the operation of software platforms dedicated to VOs (VO Management Services, user identity provisioning, VO grid services etc.), and the operation of tools to assist VO administration and monitoring. The collaboration between the active User Communities and the Resource Providers of EGI has been strengthened in PY3.

  The overall quantity of computing resources used in PY3 amounts to 12.01 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY2 amounted to 10.5 Billion HEP-SPEC 06 Hours) as shown in Table 9. The PY3 workload was generated by 507.2 Million jobs, which amounts to an average of 1.43 Million job/day.

  The overall compute resource utilization during PY3 has been significantly increasing both in terms of the cumulative number of jobs successfully done and the normalized CPU wall time consumed by all disciplines. In the refenrece period April 2012-March 2013 the rate of jobs succssfully executed increased by +8.0%, while the total normalized CPU wall time (HEP-SEPC06) incrased by +45.8%.

  While the HEP utilization is dominating in absolute terms (93.78% of the total EGI consumption), a number of other communities significantly increased their CPU wall time

[39] Contributions from other project activities to the accomplishment of the project objjectives are documented in the annual review deliverables specific to each activity.

[40] http://www.egi.eu/community/resource-providers/index.html

utilization: Earth Sciences (+123.45% yearly increase), Computational Chemistry (+78.31%), Astronomy Astro-particle and Astrophysics (+76.64%), Life Science (+65.12) and other sciences (+199.45%). Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.82% of the overall EGI used CPU wall clock time. Life Sciences are the third community for usage (1.52% of the overall EGI used normalized CPU time).

- Objective 4 (O4): *Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects*.

  EGI is actively collaborating with various ESFRI cluster projects to investigate and demonstrate the reuse of EGI core operational and infrastructural services to meet common ESFRI requirements. A collaboration was established with the EUDAT and PRACE infrastructures and user communities started in November 2012[41] aiming for the integration of data access and processing across the three infrastructures. Use cases are being collected for data access, transfer, replication and processing in various disciplines: (seismology, earth science, human physiology and hydrometeorology). Common data access and transfer tools and protocols that can be provided by all three e-infrastructures will be identified.

- Objective 5 (O5): *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*.

  The “Resource infrastructure Provider Operational Service Agreement” [RPO] was introduced in October 2011 to facilitate the exchange of operational services and the integration between the EGI-InSPIRE infrastructure and those operated by internal and external partners.

  The EGI Core Infrastructure Platform service levels were defined in the EGI.eu Operational Level Agreement [EGIO], which was approved for the first time in January 2013. This agreement is the foundation for the provisioning of operations tools as a service to other resource infrastructures.

  The EGI service registry (GOCDB) was adopted by EUDAT to support operations, and EGI-InSPIRE supported the implementation of EUDAT requirements through JRA1 development activities. EGI is currently responsible of the technical installation of the service. PRACE expressed interest in GOCDB. The version to be released in PQ13 will be tested and verified.

  A collaboration with EUDAT will be established on the evaluation of the EGI Service Availability Monitoring and its suitability to EUDAT deployment needs.

  A collaboration was also established in PQ09 with XSEDE, a major research infrastructure providing HPC resources in US. A submission of Collaborative Use Examples (CUEs) for collaborating research teams utilizing resources in EGI and XSEDE (which includes resources provided by the Open Science Grid) was opened in PQ10 with the aim of getting a better understanding of the breadth of research activities and of the usage modalities that would benefit from a XSEDE and EGI collaboration.

- Objective 6 (O6): *Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids, etc.) and heterogeneous resources(e.g.*

[41]EGI, EUDAT and PRACE workshop on data management:

https://indico.egi.eu/indico/conferenceTimeTable.py?confId=1228#20121126

HTC and HPC) into a seamless production infrastructure as they mature and demonstrate value to the EGI community.

The integration scenarios and processes of the EGI Core Infrastructure Platform supporting integrated operations of e-Infrastructure were completed and are documented in Deliverable D4.6 [D4.6].

Integration of ARC-CE, UNICORE, GLOBUS, Desktop Grid and QosCosGrid software is now complete, with the only exception of accounting whose progress was put on hold waiting for the publishing of a new accounting publisher (APEL) based on a new publishing protocol (Stomp Secure Messaging v2). This publisher was released by EMI in PQ12 and is currently under verification by EGI. All these software stacks are already deployed in production by various NGIs.

Information about the SA1 future work plan is provided by the EGI Technical Roadmap for PY4 [D2.33].

12 REFERENCES

AVL	EGI availability and reliability (https://www.egi.eu/infrastructure/Figures_and_utilisation/Availability_reliability.html)
D4.6	EGI Operations Architecture: Infrastructure Platform and Collaboration Platform Integration, EGI-InSPIRE Deliverable D4.6, Nov 2012 (https://documents.egi.eu/document/1309)
D2.33	EGI Technical Roadmap, EGI-InSPIRE Deliverable D2.33 (https://documents.egi.eu/document/1706)
EGIO	EGI.eu Operational Level Agreement, January 2013 (https://documents.egi.eu/document/1093)
GLO	EGI Glossary (https://wiki.egi.eu/wiki/Glossary)
HS06	https://wiki.egi.eu/wiki/HEP_SPEC06
MAN05	Top-BDII High Availability, EGI Manual MAN05 (https://wiki.egi.eu/wiki/MAN05)
MS512	Software Provisioning Process, EGI-InSPIRE Milestone MS512, July 2012 (https://documents.egi.eu/document/1135)
RCO	Resource Centre Operational Level Agreement v1.1, March 2011 (https://documents.egi.eu/document/31)
RPO	Resource infrastructure Provider V1.1 (https://documents.egi.eu/document/463)
SAG	South African Grid Initiative – SAGrid (https://documents.egi.eu/document/495)
SAM	Service Availability Monitoring (https://wiki.egi.eu/wiki/SAM)
SAMV	Service Availability Monitoring for VOs (https://wiki.egi.eu/wiki/Services_and_Tools_Portfolio)
SOSP	Service Operations Security Policy (https://documents.egi.eu/document/669)
SRW	Staged Rollout: https://wiki.egi.eu/wiki/Staged-Rollout
UMD	Unified Middleware Distribution (http://repository.egi.eu/category/umd_releases/distribution/umd_1/)
UNG	MoU between EGI.eu and BCC – Ukraine (https://documents.egi.eu/document/8560)
URT	UMD Release Team Terms of Reference (https://documents.egi.eu/document/1618)

13 ANNEX I. VO DISTRIBUTION PER DISCIPLINE

13.1 Astronomy Astrophysics and Astro-particle Physics

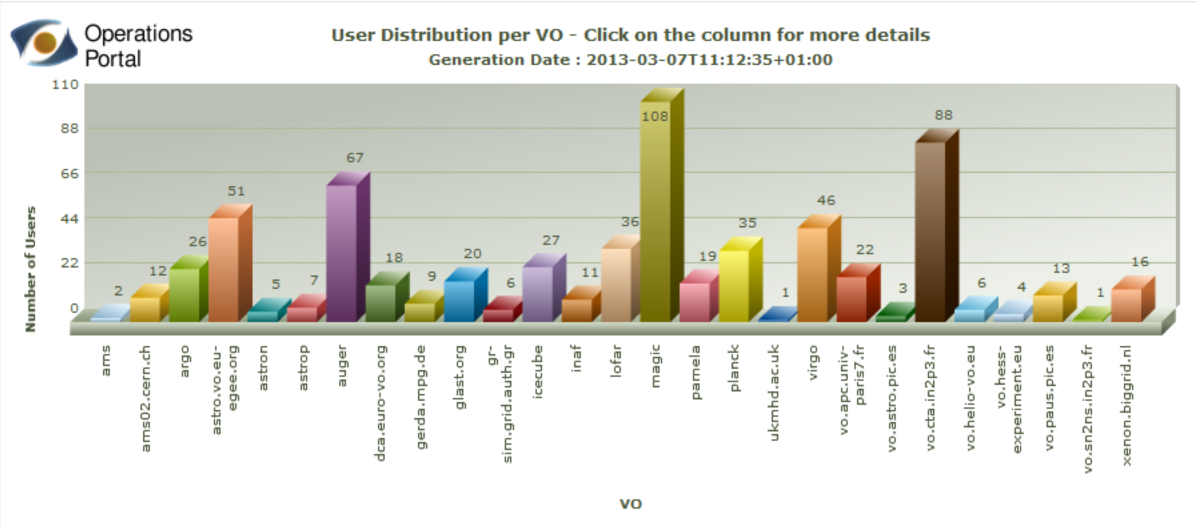

13.2 Computer Science and Mathematics

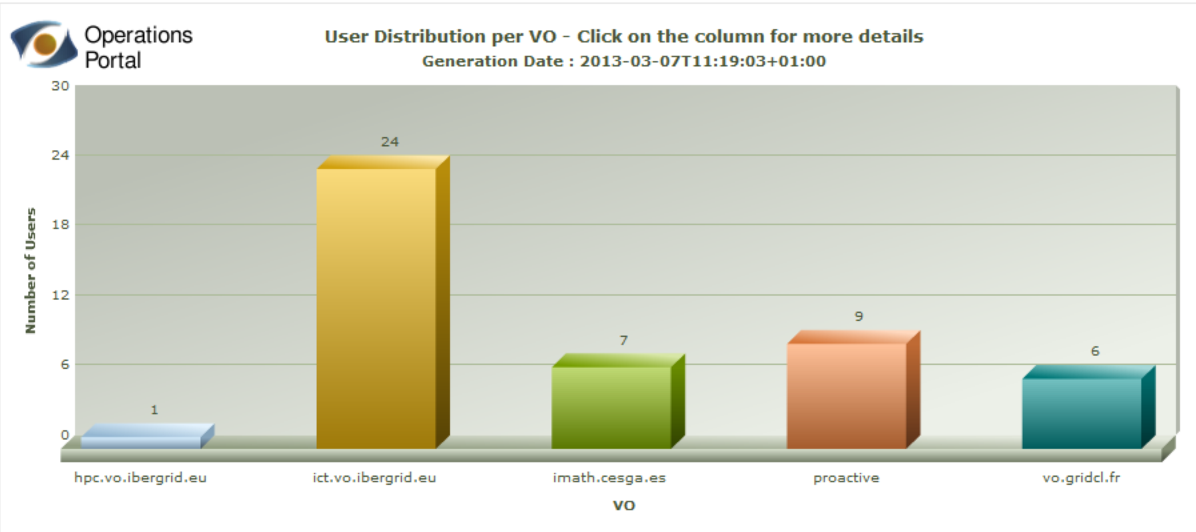

13.3 Computational Chemistry

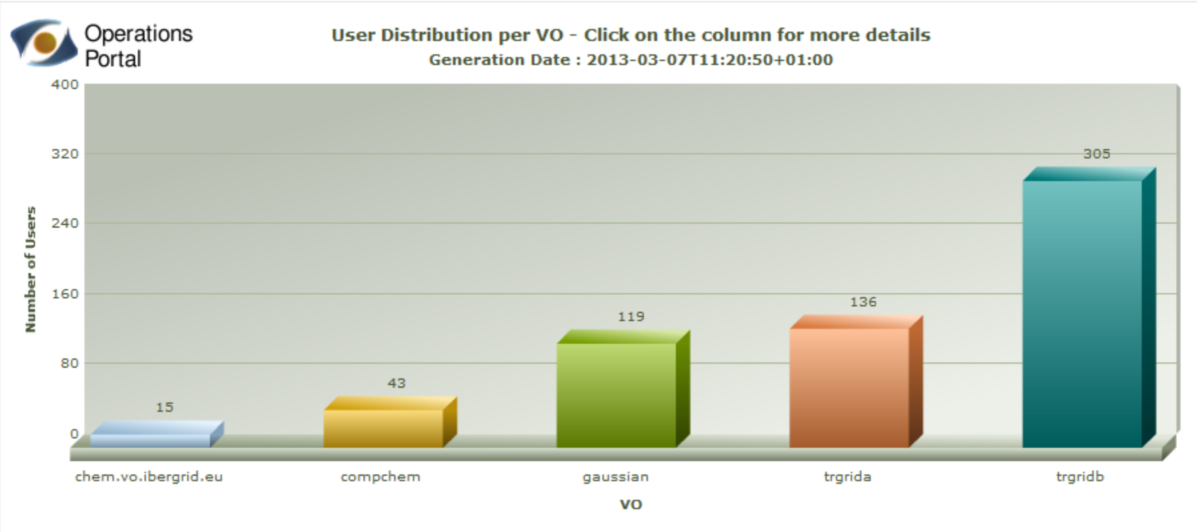

13.4 Earth Sciences

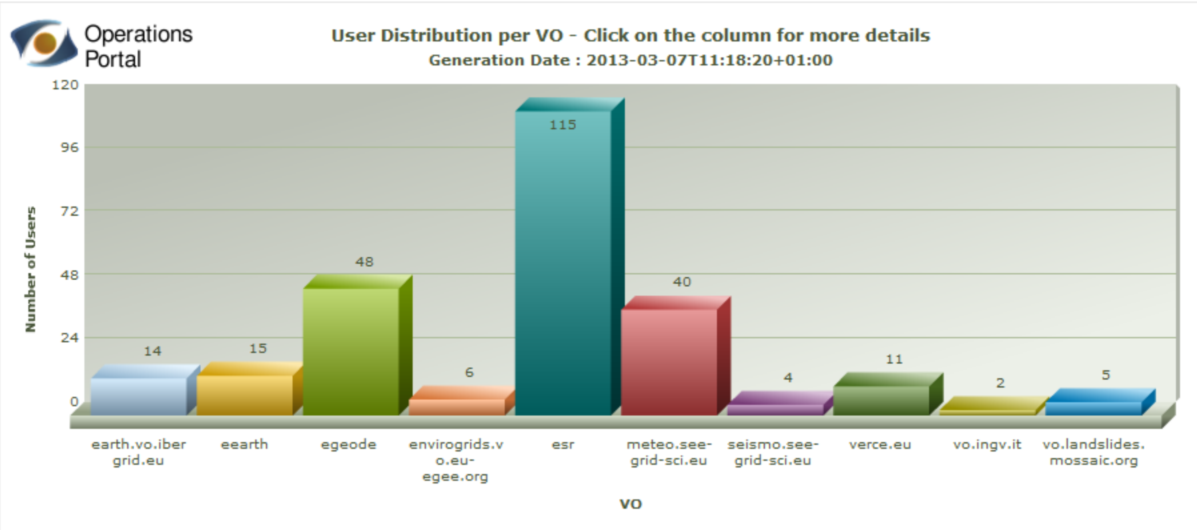

13.5 Fusion

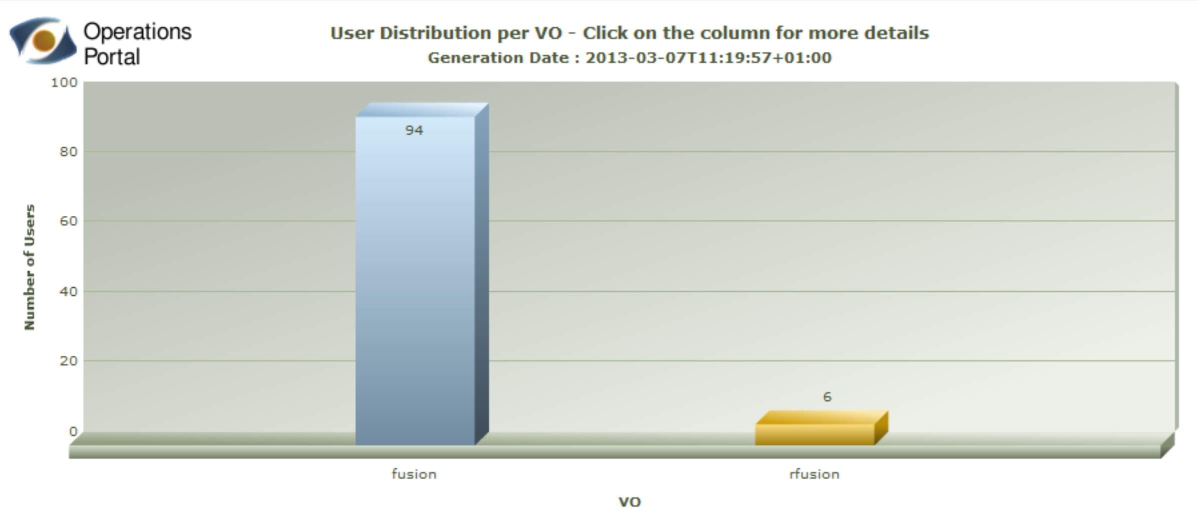

13.6 High Energy Physics

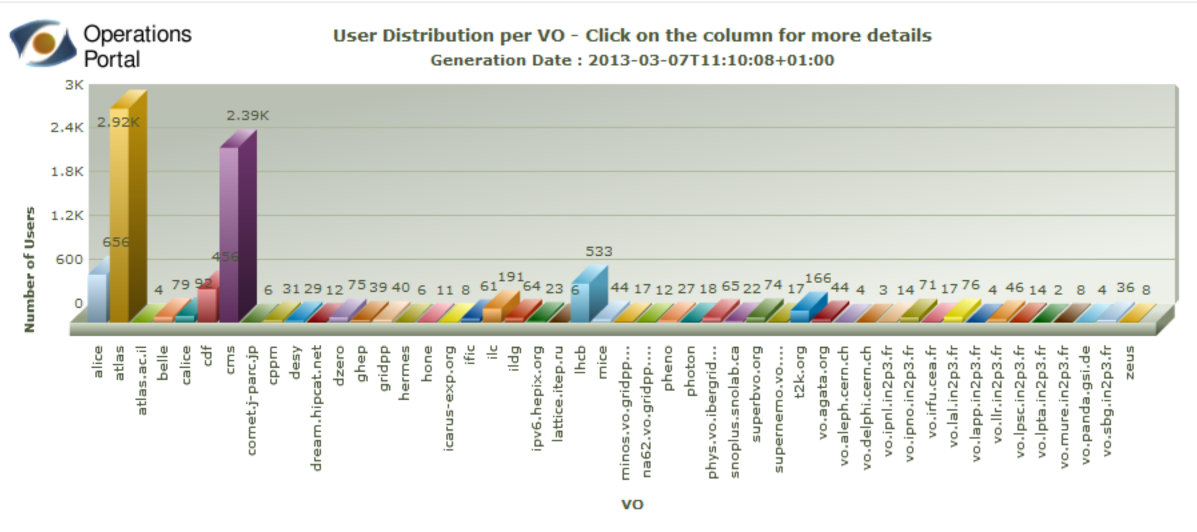

13.7 Infrastructure

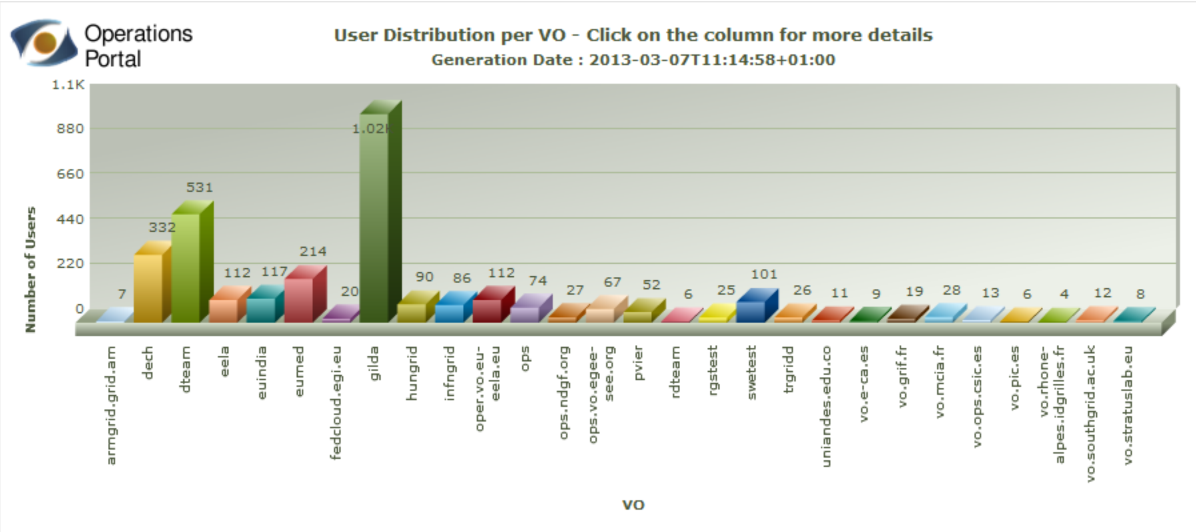

13.8 Life Sciences

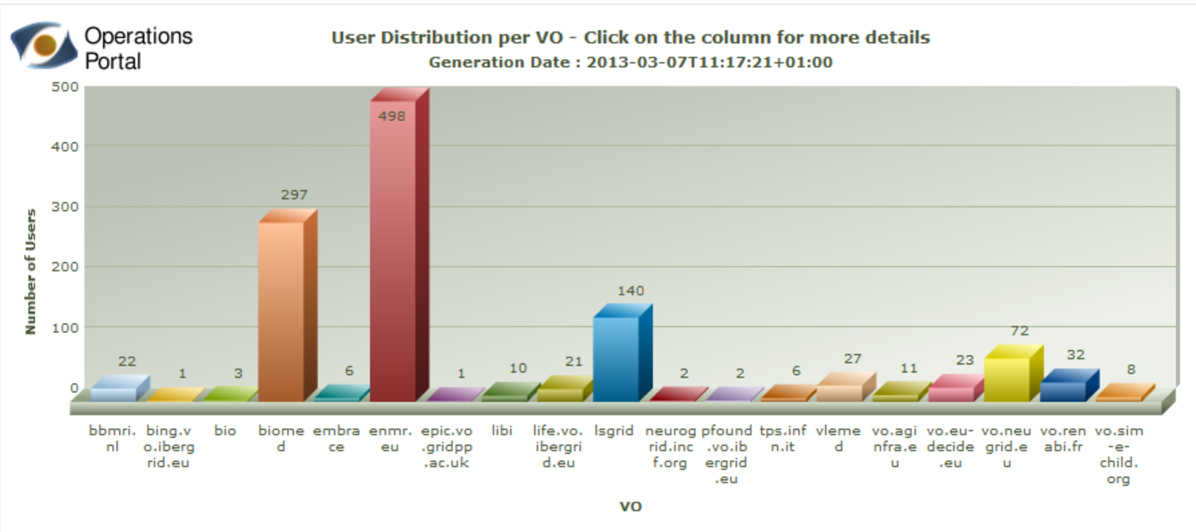

13.9 Multidisciplinary VOs

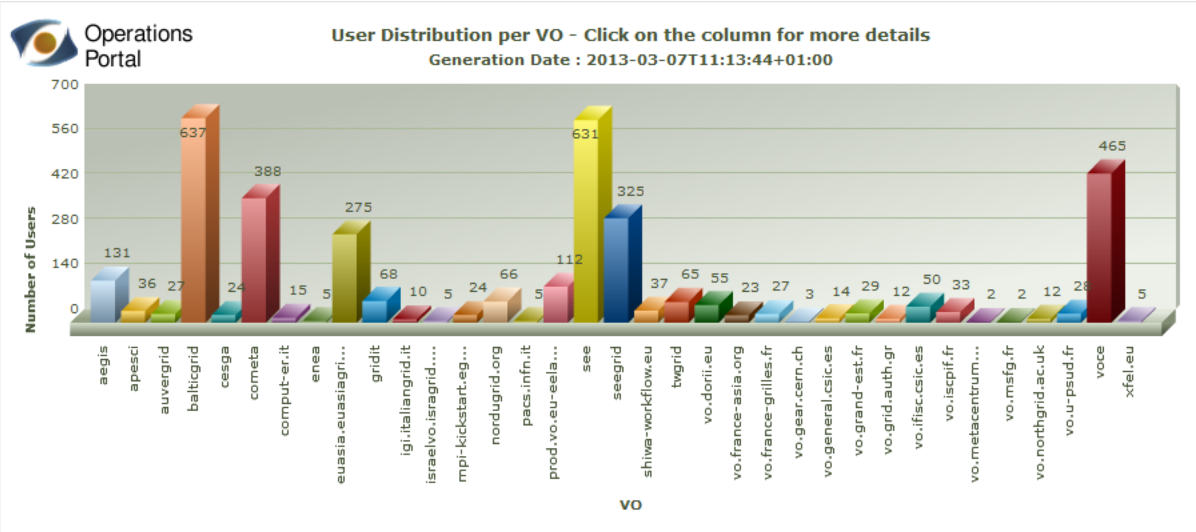

13.10 Other Disciplines