EGI-InSPIRE

ANNUAL REPORT ON THE EGI PRODUCTION INFRASTRUCTURE

EU DELIVERABLE: D4.10

Document identifier:	EGI-InSPIRE-D4.10-final.docx
Date:	**18/06/2014**
Activity:	**SA1**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/2240

Abstract

This document provides information on the status of the EGI Resource Infrastructure at the end of PY4. 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.

In compliance with the Third Amendment of the project, this deliverable also reviews the status of the operations Global Tasks (formerly MS130) and NGI International Tasks (formerly MS131).

I. DELIVERY SLIP

	Name	Partner/Activity	Date
From	M. Krakowian	EGI.eu/SA1	28-05-2014
Reviewed by	**Moderator:** Tiziana Ferrari **Reviewers:** Peter Solagna	EGI.eu/NA1 EGI.eu/SA2	10-06-2014
Approved by	**AMB & PMB**		17-06-2014

II. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	28-05-2014	First draft for external review and AMB review	M. Krakowian/EGI.eu

III. 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.

IV. DOCUMENT AMENDMENT PROCEDURE

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

V. TERMINOLOGY

A complete project glossary is provided at the following page: https://wiki.egi.eu/wiki/Glossary.

VI. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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

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

VII. EXECUTIVE SUMMARY

During PY4 SA1 was responsible of the continued operation and expansion of the production infrastructure. The total number of Resource Centres (RCs) in March 2014 amounts to 355 certified instances (+3% yearly increase), and it includes one Cloud SME. The increase in the number of resource centres has been mainly driven by the integration of new resource infrastructure providers in the Africa-Arabia region as an outcome of the collaboration agreement signed in February 2014[1], and towards the end of PY4 to the preparatory work leading to the launch of the EGI Federated Cloud infrastructure, which required the certification of new Cloud providers contributing services to the Federated Cloud Testbed, and whose operations integration is the main outcome of PY4 activities in SA1. In PY4 four new countries joined the infrastructure: Colombia, Mexico, Morocco and South Africa.

The offered capacity substantially increased: +20% yearly increase for compute resources in PY4, while disk capacity increased to 250 PB (+29% yearly increase). In March 2014 the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to about 434,000, which provide 3.5 Million HEP-SPEC 06.

The use of gateways to provide users with a native user-friendly environment to the infrastructure services is increasing. Quite often user portals provide users with the capability of using institutional credentials to authenticate themselves, these credentials are then mapped to robot certificates (often owned by the VO managers). By doing so registration of users to a VO and obtaining personal X.509 certificates are not necessary, this contributes to increase the user friendliness of the platforms. Use of robot certificates is internally accounted for by the portals in compliance to the VO Portal policy[2]. In PY4 the number of robot certificates embedded in user gateways is 114; robot certificates are used by 40 VOs in total. 15,000 users can potentially use scientific gateways; this is increased by the number of registered users to active VOs, which is estimated to be 20,478 at the end of PY4.

The overall quantity of computing resources *used* in PY4 amounts to 15.1 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY3 amounted to 12.01 Billion HEP-SPEC 06 Hours). The PY3 workload was generated by 587.8 Million jobs, which amounts to an average of 1.61 Million job/day.

The overall compute resource utilization during PY4 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 reference period April 2013-March 2014 the rate of jobs succssfully executed incrased by +10.8%, while the total normalized CPU wall time (HEP-SEPC06) incrased by +15.98%.

The High Energy Physics is the largest community in terms utilization, however during PY4 it decreased by -1.5% when comparing it to PY3. HEP utilization amount to 88.34% of the total EGI resourced utilization in normalized CPU wall time hours. A few other disciplines - which gather users from international research collaborations and the long tail – significantly increased their relative utilization in comparison to PY4 (CPU wall time utilization): Fusion (+108.46% yearly increase), Computational Chemistry (+48.31%), Multidisciplinary (+76.13), Infrastructure (+14.75%) and unknown disciplines (+6417.66%). As of PY5 a revised VO categorization that classifies projects into

---

[1] The Memorandum of Understanding involved the following countries: Algeria, Egypt, Ethiopia, Ghana, Nigeria, Kenya, Morocco, Senegal, South Africa (leading partner), Tanzania and Tunisia.

[2] https://documents.egi.eu/document/80

44 different scientific disciplines has been adopted[3]. The new classification allows a fine-grained grouping of users and accounting information, which provide a better mechanism to track trends and impact of user engagement activities. The operational tools (mainly the VO registration system and the accounting portal) are being adapted to it.

The performance of NGIs was tracked in terms of NGI core middleware services and NGI core operations tools in order to report on the workload and the quality of the services provided at a national level. For NGI core middleware services Availability and Reliability were 98.7% and 99.3% respectively, when for NGI core operations tools 98.21% and 98.68%.

Monitoring of the 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 – and since PY4 have been divided in to two categories: *critical* central operations tools and *non-critical* central operations tools. Availability and Reliability were 97.95% and 97.95% respectively for critical central operations tools, 99.47% and 99.47% respectively for non-critical central operations tools.

The EGI infrastructure as a whole includes a total number of 3300 individual service interfaces, including infrastructural services (job submission, workload management, data management services etc.) as well as community services (gateways, portals, VO-specific platforms etc.). Various grid middleware stacks are in production in EGI. An indication of their distribution is given by the various Compute Element services deployed by Resource Centres. CREAM-CE is in production in the 88.89% of the infrastructure, ARC-CE is second in deployment (10.56%) followed by QCG.Computing (1.82%), Unicore6.TargetSystemFactory (1.28%) and GRAM (0.55).

In PY4 cloud resources have been introduced in production by certifying the 13 first cloud providers. The same set of core infrastructure services is being used to operate the federation of grid and cloud services, this meaning that no ad-hoc procedures and tools are necessary to operate a pan-European standard based cloud infrastructure.

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

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 PY4. 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 9. Section 10 provides information about Software support Unit. Global Operations Tasks review is described in Section 11. Section 12 summarises the progress made in operations during PY4 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 2014).

Resource Centres	Number of RCs (certified)
EGI-InSPIRE Partners and NGI Council Members/PY4 Target	316 (does not include 23 suspended RCs)
From non-European EGI-InSPIRE Partners	24
From integrated Infrastructures (Canada, Latin America)	39
Total	**355/378**

As shown in the table above, the total number of certified RCs in March 2014 amounts to 355 instances, of which: 316 are contributed by European NGIs/EIROs that are EGI-InSPIRE partners or Council members and 39 by integrated RPs namely: Canada, Latin America – Brazil, Chile, Argentina, Venezuela and Mexico.

Of the 316 RCs mentioned above, 27 are contributed by Asia Pacific NGIs.

The project target for PY4 of 345 RCs was met (355 RCs certified centres against 345 expected), exceeding the target for a 3% deviation.

The increase in the number of resource centres has been mainly driven by the integration of new resource infrastructure providers, such as Africa-Arabia, and towards the end of PY4 the certification of new sites providing exclusively cloud resources. In PY5 the plan is to extend the amount of cloud resources and sites from integration new resource infrastructures providers: IHEP-China, C-DAC India.

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

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

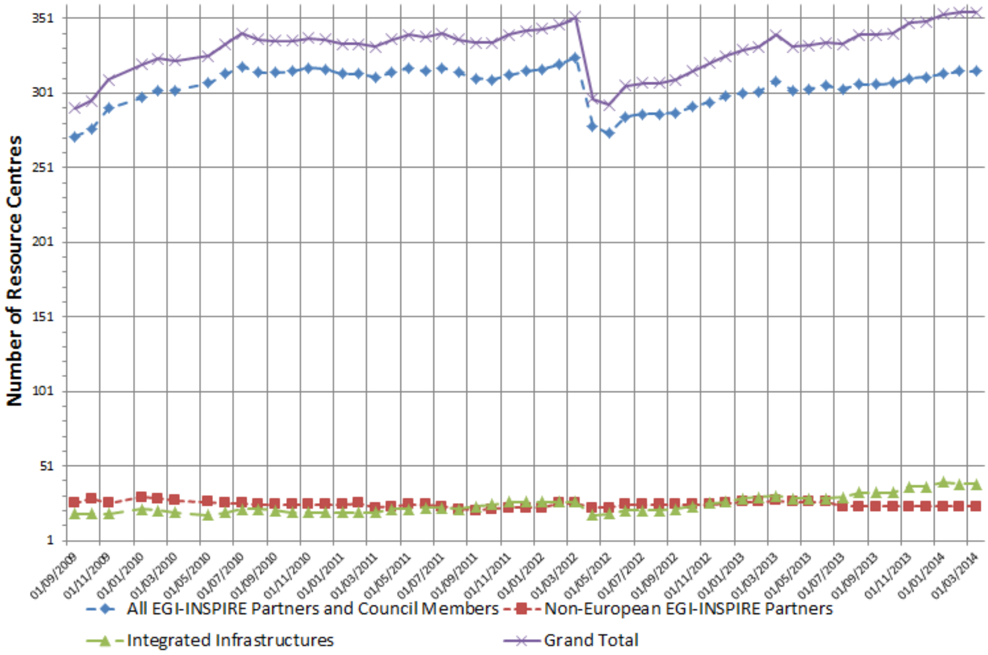

Figure 1. Number of certified production RCs from 01/09/2009 to 01/03/2013. (data source: GOCDB).

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 2014 EGI comprises resources provided across 54 countries and one European Intergovernmental Research Institute – CERN.

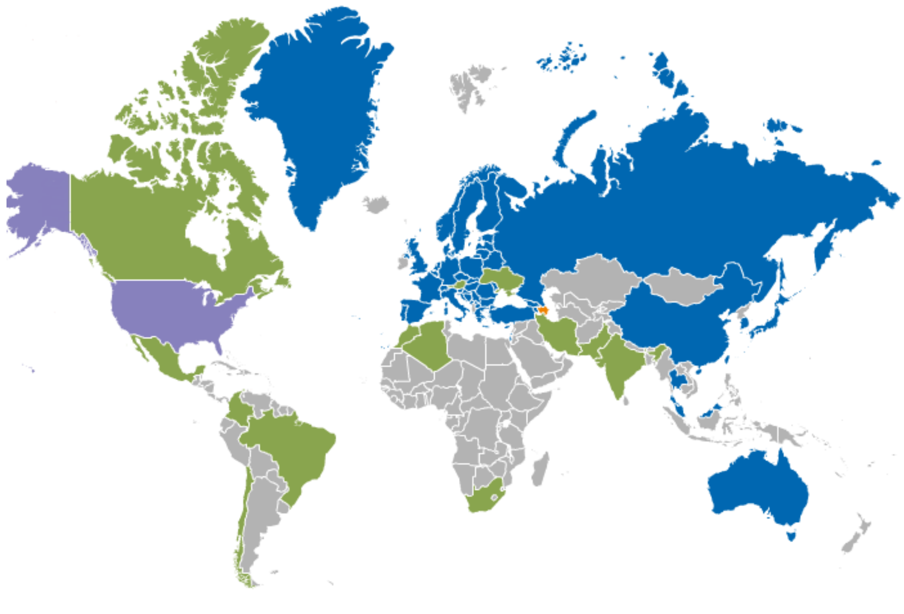

Figure 2 RPs distribution in March 2014 (data source: GOCDB). Legend: (blue) Integrated EGI-InSPIRE Partners and EGI Council Members, (green) External Resource Providers, (orange) Internal/External Resource Providers – Azerbaijan, (purple) Peer Resource Providers – Open Science Grid.

EGI-InSPIRE partners or RPs that are Council members contribute resources from 44 countries. However, 10 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, Iran (Asia Pacific Federation);
- Austria (Italian NGI): operations support for Austrian RCs is provided by the NGI_IT Operations Centre;
- Canada (Canada NGI);
- Brazil, Chile, Colombia, Mexico, Argentina (Latin American Federation)
- South Africa (Africa-Arabia Federation)

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

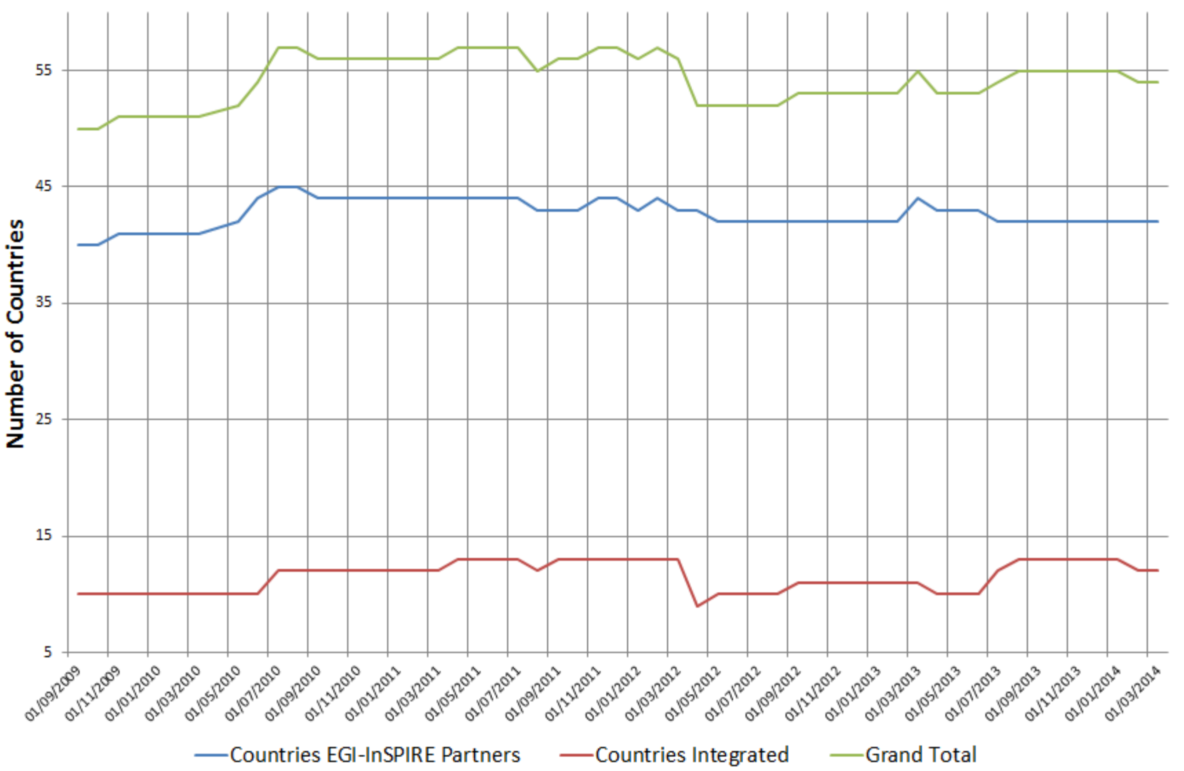

Figure 3. EGI countries hosting certified production Resource Centres from 01/09/2009 to 01/03/2014 (data source: GOCDB).

The number of countries contributing resources is approximately constant. In PY4 four countries stopped contributing resources: Belarus, Philippines, Vietnam and New Zealand; while three countries started contributing resources: Colombia, Moldavia and South Africa. The trend since 01 September 2009 is illustrated in Figure 3.

In PY4 the production infrastructure in Belarus has been moving to new middleware, and RCs have been periodically temporarily removed from the infrastructure; this work will be finished in PY5, and the NGI plans to contribute again to the EGI infrastructure by the end of 2014.

The integration of the South African infrastructure, and in the future the other countries that will be part of the Africa Arabia federation, expanded the EGI collaborations in new continent. Currently EGI has integrated partners in all the five continents, laying down the basis for future collaborations with research projects and user communities based also in countries outside Europe who need access to computing and storage resources.

Table 2. Distribution of production certified RCs across countries and Operations Centres in PQ12, PQ13, PQ14, PQ15 and March 2014 (data is sorted by country). (Data source: GOCDB).
I = Integrated, EP = European Partner, P = non-European partner, C = only Council member

Operations Centre	Country		PQ12	PQ13	PQ14	PQ15	01/03/2014
Latin America	Argentina	I	0	1	1	1	0
NGI_ARMGRID	Armenia	EP	2	2	2	1	1
Asia Pacific	Australia	P	1	1	1	1	1
NGI_IT	Austria	I	2	2	2	2	2
NGI_BY	Belarus	EP	1	0	0	0	0
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	3	3	3	5	5
Canada	Canada	I	8	8	8	8	8
Latin America	Chile	I	1	1	1	2	2
Asia Pacific/Canada	China	P	2	2	2	2	2
Latin America	Colombia	I	0	1	1	1	1
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	3	3
NGI_NDGF	Denmark	EP	2	2	2	2	2
NGI_NDGF	Estonia	C	2	2	2	2	2
NGI_FI	Finland	EP	11	11	11	11	11
NGI_FRANCE	France	EP	18	19	19	19	19
NGI_MARGI	FYROM	EP	3	3	3	3	3
NGI_GE	Georgia	EP	1	1	1	1	1
NGI_DE	Germany	EP	19	19	21	20	20
NGI_GRNET	Greece	EP	16	16	16	16	16
NGI_HU	Hungary	EP	3	4	4	4	4
Asia Pacific	India	I	2	1	1	2	2
Asia Pacific	Iran	I	1	1	1	1	1
NGI_IE	Ireland	EP	0	0	0	0	0
NGI_IL	Israel	EP	7	8	8	8	8
NGI_IT	Italy	EP	53	51	52	53	55
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	3	3	3	3
Latin America	Mexico	I	1	1	1	1	1
NGI MD	Moldova	EP	0	3	3	4	4
NGI_ME	Montenegro	EP	1	1	1	1	1
NGI_NL	Netherlands	EP	16	15	15	14	14
NGI_NDGF	Norway	EP	1	1	1	1	1
Asia Pacific	Pakistan	P	2	2	2	2	2
NGI_PL	Poland	EP	10	10	10	11	11
NGI_IberGrid	Portugal	EP	6	7	7	7	7
NGI_RO	Romania	EP	10	10	10	10	10
Russia	Russia	EP	11	11	11	11	11
NGI_AEGIS	Serbia	EP	5	6	6	6	6
NGI_SK	Slovakia	EP	4	4	4	7	7
NGI_SI	Slovenia	EP	2	2	2	2	2
South Africa	South Africa	I	0	0	3	6	6
Asia Pacific	South Korea	P	5	4	4	5	5
NGI_IberGrid	Spain	EP	21	21	21	20	20
NGI_NDGF	Sweden	EP	1	1	1	2	2
NGI_CH/CERN	Switzerland	EP	8	7	8	8	8
Asia Pacific	Taiwan	P	6	6	6	6	6
Asia Pacific	Thailand	P	3	3	3	2	2
NGI_TR	Turkey	EP	3	3	3	3	3
NGI_UA	Ukraine	I	12	12	12	14	14
NGI_UK	United Kingdom	EP	23	24	24	24	24
Asia Pacific	Vietnam	P	1	0	0	0	0
TOT Resource Centres			**332**	**334**	**341**	**354**	**355**
TOT Countries			**53**	**54**	**55**	**55**	**54**

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 30 national operations centres and 8 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, where the amount of sites per country do not justify the creation of a national operations centre, but suggests a international collaboration. The creation of new national grid initiatives in those regions will depend on their expansion plans and on national policies.

In PY4 new Operations Centre for South Africa (NGI_ZA) has been created and due to decision to create federated Operations Centre for Africa and Arabia it will be decommission in PY5. All sites have been already moved to newly created Africa-Arabia Operations Centre.

Two new Operations Centres are in phase of creation:

- NGI_China: will cover all Chinese sites, currently registered under the Canada NGI
- IDGF: will cover sites which provide Desktop Grid solution through IDGF project

Federated operations centres	Member countries	Comments
Asia Pacific	Australia, China, India, Japan, Malaysia, Pakistan, South Korea, Taiwan, Thailand, Iran	China sites will be moved to new NGI under creation to capture all Chinese sites
Canada	Canada, China	China sites will be moved to new NGI under creation to capture all Chinese sites
IberGrid	Portugal, Spain	
Latin America (ROC_LA)	Brazil, Chile, Colombia, Mexico, Argentina, Brazil, Chile,	
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	
Arabia Africa	South Africa	Operations Centre will cover sites from Africa and Arabia

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, GLOBUS and Cloud 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, Cloud and UNICORE RCs.

3.1 Compute Resources

At the end of PQ15 the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to 401,591, which provide 3.5 Million HEP-SPEC 06, while the total number including compute resources contributed by integrated and peer infrastructures amounts to 532,237 units. This value exceeds the PY4 target of 400,000 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 7.34% since March 2013, while the installed capacity in HEP-SPEC 06 increased by 1.03%.

Table 4. EGI-InSPIRE logical CPUs

Logical CPUs	PQ15/PY4 Target
EGI Council participants	401,591/400,000
EGI-InSPIRE partners plus integrated and peer infrastructures	532,237

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, March 2013 and March 2014

	April 2010 (EGEE-III Infrastructure)	March 2011 (EGI-InSPIRE Infrastructure)	March 2012 (EGI-InSPIRE Infrastructure)	March 2013 (EGI-InSPIRE Infrastructure)	March 2014 (EGI-InSPIRE Infrastructure)
Logical CPUs/yearly increase	192,000	207,203/+7.9%	270,800/+30.7%	373,235/+33.6%	433,957/+20 %

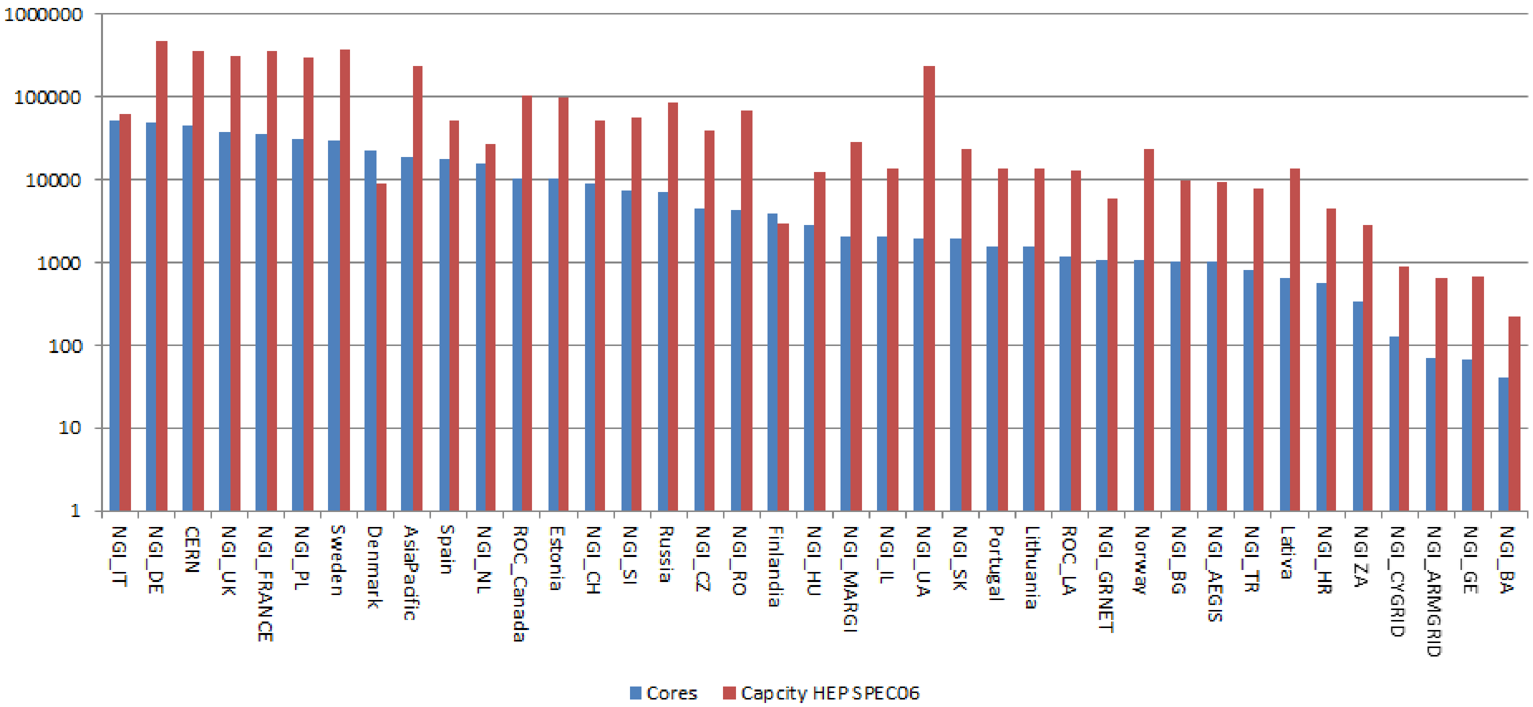

Figure 4. Log scale distribution of logical cores (blue bar) and HEP SPEC 06 installed capacity (red bar) at the end of PQ15 across EGI Resource infrastructure Providers, including EGI-InSPIRE partners, EGI Council members and integrated infrastructures. Data sorted by number of cores. Source: project 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 PQ15, the total amount of reported installed disk capacity amounts to 250 PB (+29% yearly increase). The distribution of disk storage resources among the EGI-InSPIRE partners is illustrated in Figure 5, which shows that disk capacity is concentrated across six NGIs/EIROs, which are in descending order (for PQ15): Italy, USA, Germany, CERN, Asia Pacific and United Kingdome.

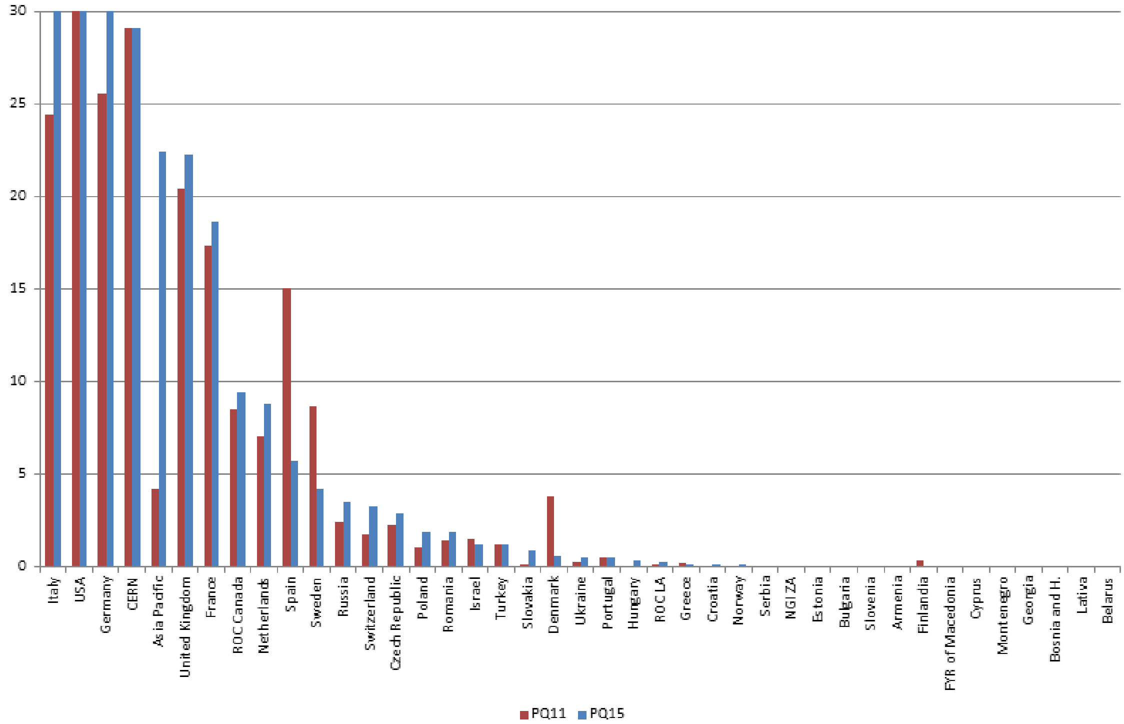

Figure 5. Installed disk capacity in PB across the EGI RPs at the end of PQ11 – red bar – compared to the installed capacity in PQ15 – blue bar (source: Metrics Portal and Gstat).

Tape capacity is mainly provided by CERN and WLCG Tier-1 RCs. At the end of PQ15 the total installed tape (also called *nearline storage*) capacity reported in Gstat amounts to 168.8 PB.

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 32 units (23% yearly decrease with respect to PQ11) at the end of PQ15.

Table 6. Integration metrics (HPC and MPI)

Metric	PQ11	PQ15
Number of HPC clusters (M.SA1.Integration.1)	42	32
Number of sites with MPI (M.SA1.Integration.2)	89	76

At the end of PQ15 Message Passing Interface [MPI] jobs were supported by 76 sites (-14% 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 and continue in PY4. 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 has been rolled out to production: information about MPI capabilities are not only published by services to the Information Discovery Service as today, but it is also be registered into the EGI service registration facility (GOCDB). A more accurate mechanism to estimate the MPI support in the infrastructure is available. In addition, starting with the EMI 3 release, the new APEL publisher is capable of reporting accounting information of multi-core jobs. The new publisher is gradually deployed by the production infrastructure. Accounting information of MPI jobs will be a more accurate indicator of the amount of parallel computing workload supported by EGI and are planned to be available in PY5.

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

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[8]. A new science classification was proposed in PY3 and discussed with user communities, which will adopted at the beginning of PY5.

The overall number of international and national VOs registered in the Operations Portal[9] at the end of April 2014 amounts to 211 (-9.8% from April 2013), the decrease is due to a campaign for the de-registration of expired user projects from the VO registry of EGI.

The use of gateways to provide users with a native user-friendly environment to the infrastructure services is increasing. Quite often user portals provide users with the capability of using institutional credentials to authenticate themselves, these credentials are then mapped to robot certificates (often owned by the VO managers). By doing so registration of users to a VO and obtaining personal X.509 certificates are not necessary, this contributes to increase the user friendliness of the platforms. Use of robot certificates is internally accounted for by the portals in compliance to the VO Portal policy . In PY4 the number of robot certificates embedded in user gateways is 114; robot certificates are used by 40 VOs in total. 15,000 users can potentially use scientific gateways; this is increased by the number of registered users to active VOs, which is estimated to be 20,478 at the end of PY4.

The increase in the number of Robot Certificates indicates that users, in particular new user communities, are looking for alternative authentication mechanisms different from the plain X.509 certificates. This is also signal for EGI to explore different authentication technologies in the future, or to work on a better integration of robot certificates with the production infrastructure.

The diagram below shows the trend in use of robot certificates and VOs since PY1.

[8] "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.

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

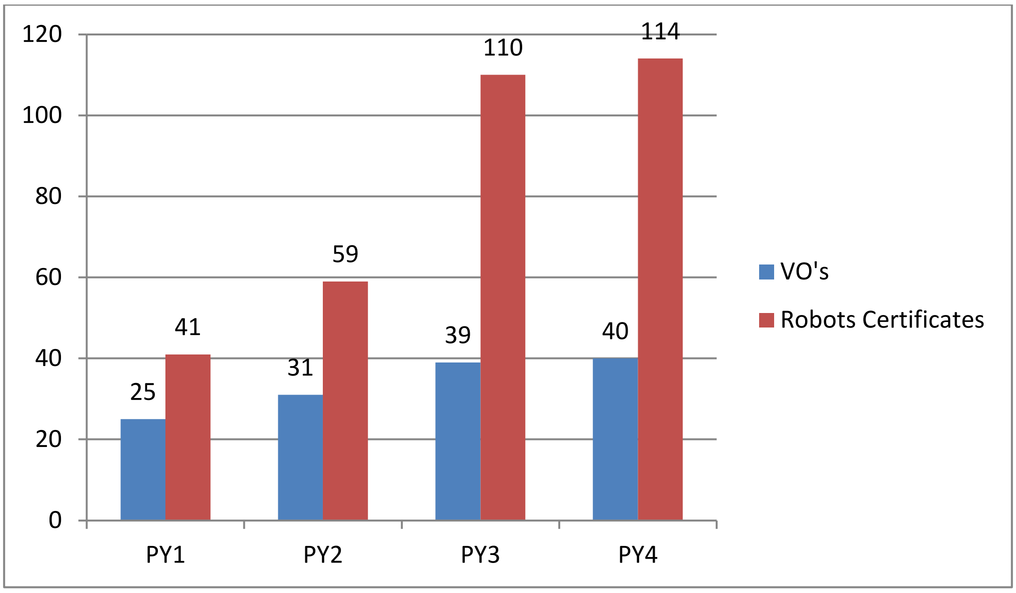

Table 7. Use of robot certificates and related VO in EGI since EGI-InSPIRE PY1.

4.1 VO Distribution across scientific fields

The distribution of VOs per discipline is illustrated in Figure 76.

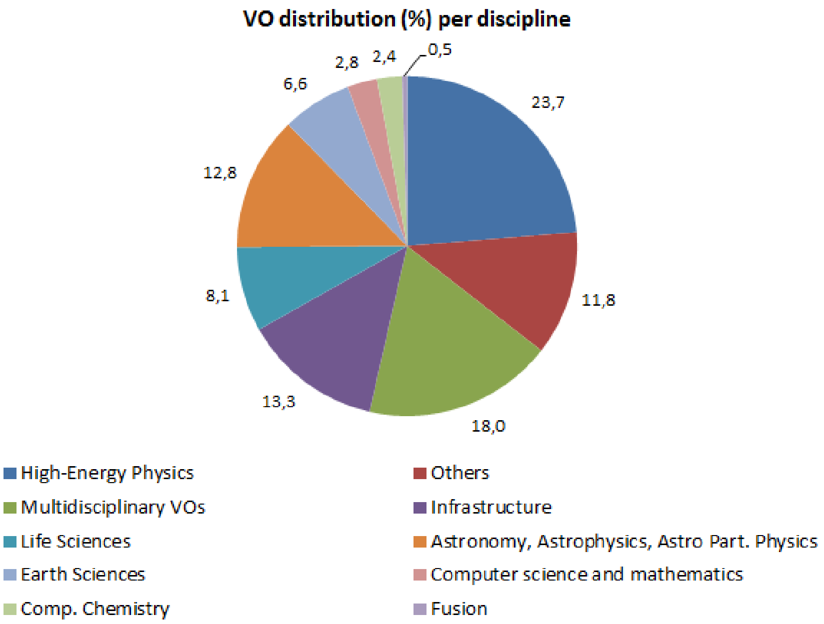

Figure 6. Distribution of number VOs per discipline (April 2014, source: Operations Portal).

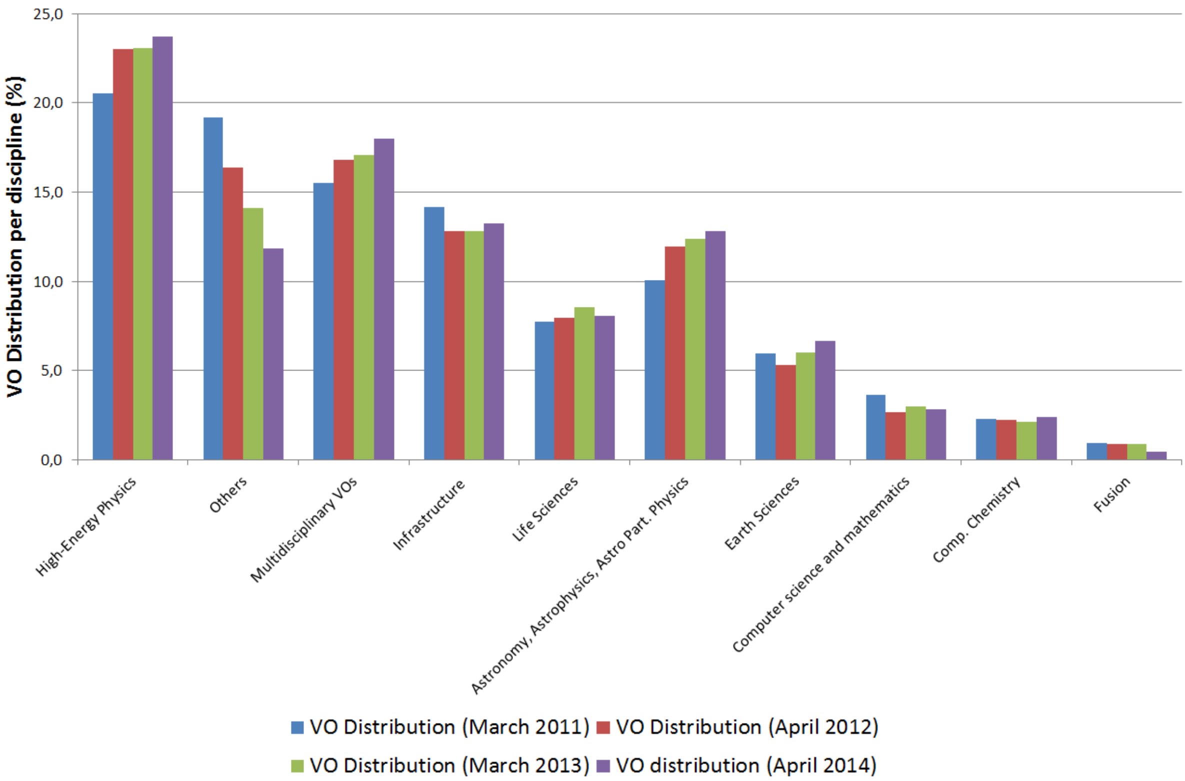

Figure 7. Comparison of the VO distribution at the end of March 2011 (blue bars), at the end of April 2012 (red bars), at the end of March 2013 (green bars) and at the end of April 2014 (purple bars). Source: Operations Portal.

4.2 User Distribution across scientific fields

The largest disciplines in terms of number of registered users are: High Energy Physics (33.24%), Infrastructure (15.69%), Multidisciplinary VOs (12.35%) and Others (12.33%). During PY4 the number of registered users for most of disciplines has increased: High-Energy Physics (+2.05) Infrastructure (+4.49%), Life Science (+15%), Astronomy Astrophysics and Astro-particle Physics (+13.56), Earth Science (+26.39%), Computer Chemistry (+4.65%), Fusion (+1.04%) and Computer Science (+15.22%). The detailed user distribution per discipline is presented in Appendix.

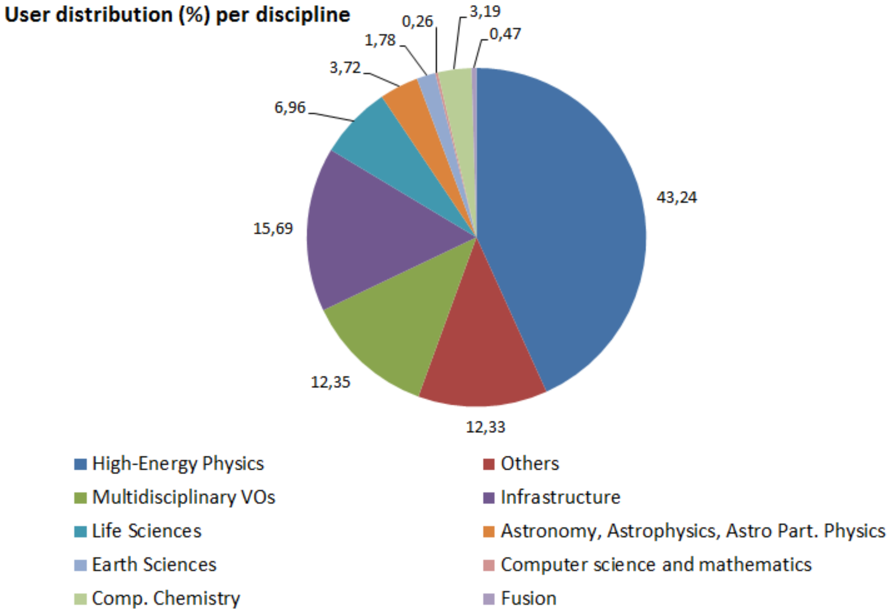

Figure 8. User distribution per discipline (April 2014, source: Operations Portal)

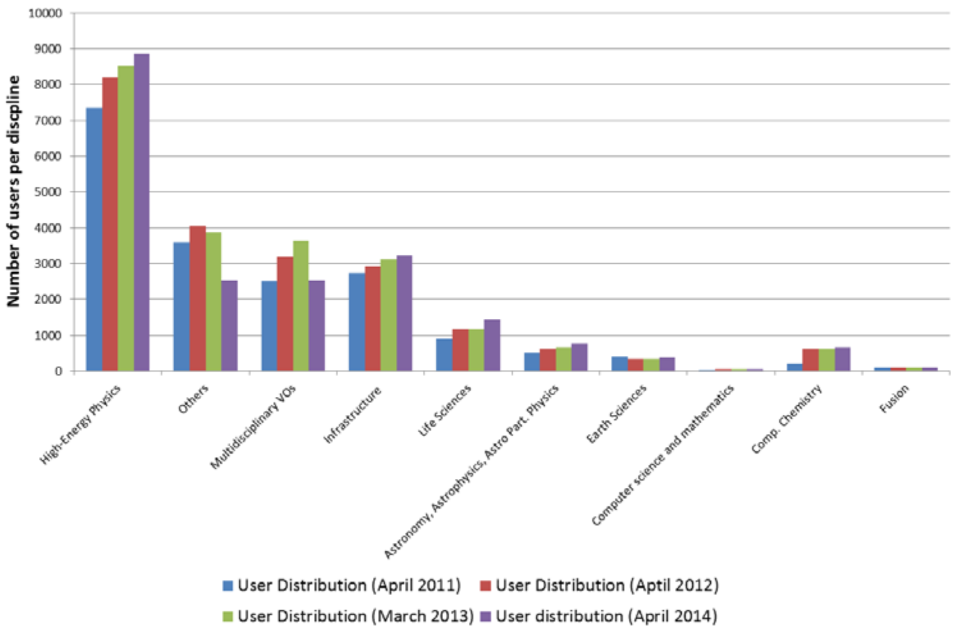

Figure 9. Comparison of the number of users per discipline in April 2011 (blue bar), April 2012 (red bar), at the end of March 2013 (green bars) and at the end of April 2014 (purple bars). Source: Operations Portal.

4.3 Resource Utilization per Discipline

Table 8. 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%
	Apr 2013 – Mar 2014	587.8	+10.8%
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%
	Apr 2013 – Mar 2014	17.84	15.98%

The overall compute resource usage during PY4 has increased 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 2013-March 2014 the rate of jobs succssfully executed incrased by +10.8%, while the total normalized CPU wall time (HEP-SPEC06) incrased by +15.98%. Table 8 compares the April 2013 – March 2014 increase trends with those achived in the previous 3 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, PY3 and PY4 are shown in Figure 10(a).

The High-Energy Physics discipline (contributing 43.24% of the user community) utilizes the highest amount of resources: 89.58% of the overall EGI amount of normalized CPU wall time hours consumed. As indicated in Table 9, the HEP usage yearly slightly decrease amounts to 1.46%.

While the HEP utilization is dominating in absolute terms, a number of other communities significantly increased their CPU wall time utilization: Fusion (+108.46% yearly increase), Computational Chemistry (+48.31%), Multidisciplinary (+76.13), Infrastructure (+14.75%) and unknown disciplines (+6417.66%).

Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.13% of the overall EGI used CPU wall clock time. Life Sciences are the third community for usage (1.80% of the overall EGI used normalized CPU time). For this community the job submission pattern shows a further increase of the job submission rate (+43.57%). This trend already started in PY2, during which the rate already showed a relative increase of +42.54%.

Life Science, Multidisciplinary science, Fusion and Computational Chemistry 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 12. 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.

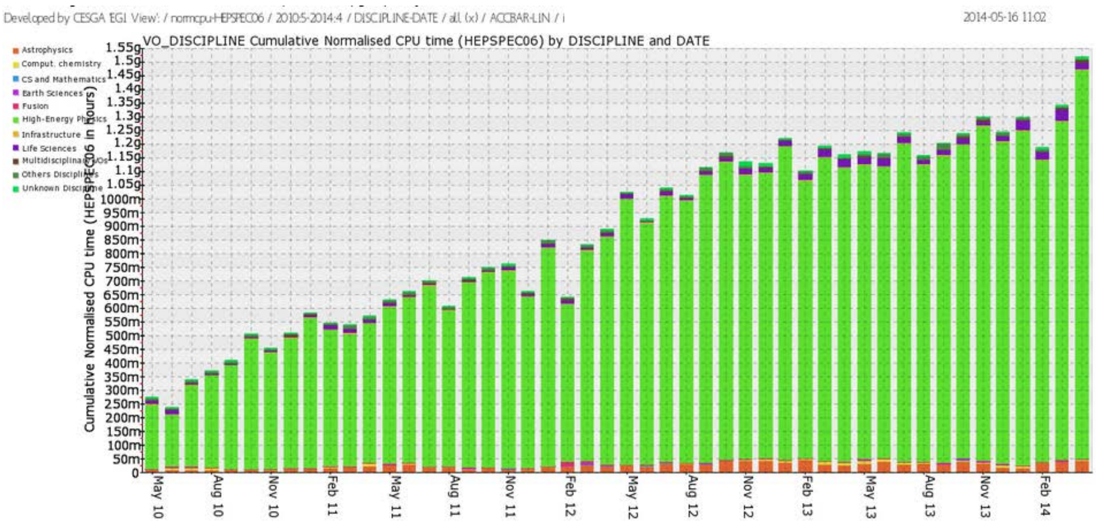

(a)

  PUBLIC

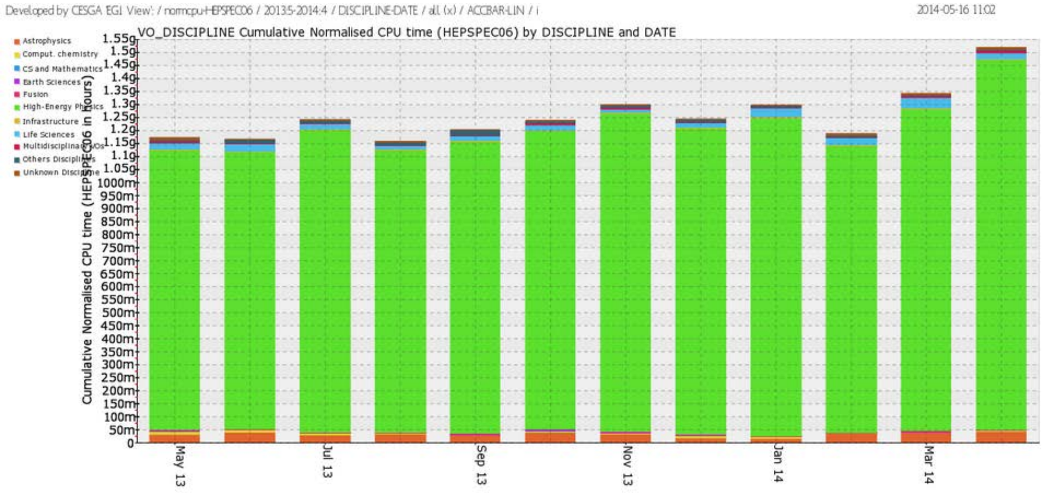

(b)

Figure 10. Usage of EGI resources (HEP-SPEC 06 CPU wall clock hours) from (a) the beginning of the project to date, and (b) during PY4 (source: accounting portal).

Table 9. Used normalized CPU wall clock time and jobs done per discipline in PY3, PY4 and the respective yearly increase (source: accounting portal).

Discipline	May 12–April 13		May 13 – April 14		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 12) (E)	(yearly increase from May 12) (F)
High-Energy Physics	93.78	89.58	88.34	76.18	**+14.38%**	-1.46%
Infrastructure	0.10	2.88	0.03	2.85	-62.72%	**+14.75%**
Life Sciences	1.52	4.34	1.80	3.78	**+43.57%**	**+0.94%**
Astrophysics	2.82	1.82	2.13	1.00	-8.09%	-36.15%
Multidisciplinary	0.12	0.17	0.28	0.26	**+183.57%**	**+76.13%**
Others Disciplines	0.59	0.45	0.43	0.24	-10.96%	-37.17%
Unknown Discipline	0.43	0.27	6.30	15.18	**+1675.94%**	**+6417.66%**
Comput. Chemistry	0.48	0.22	0.53	0.29	**+33.47%**	**+48.31%**
Fusion	0.01	0.10	0.08	0.18	**+1205.27%**	**+108.46%**
Earth Sciences	0.15	0.11	0.07	0.04	-41.35%	-54.78%
CS and Mathematics	0.00	0.07	0.00	0.00	-89.97%	-94.38%

Figure 11. Distribution of consumed normalized CPU wall time among the main active VOs (May 2013-April 2014) Figure 11 shows the distribution of consumed normalized CPU wall time among the main VOs.

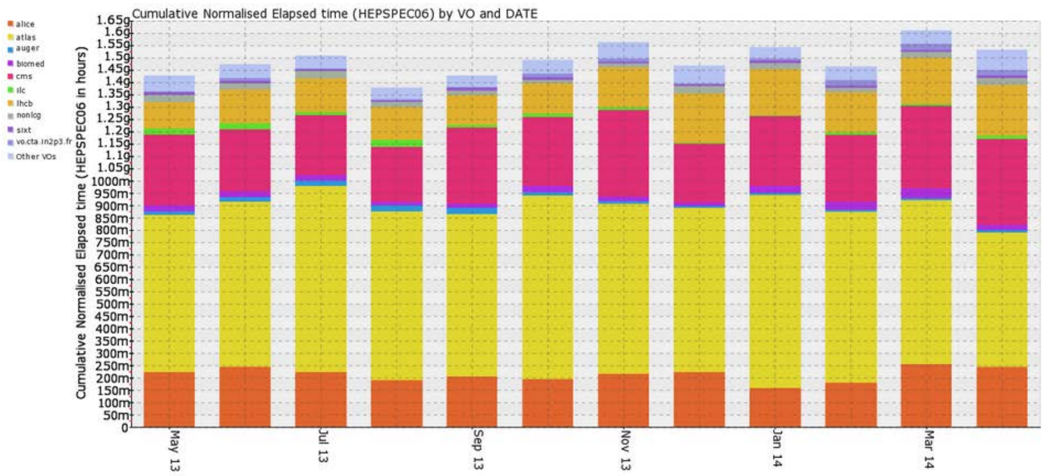

Figure 11. Distribution of consumed normalized CPU wall time among the main active VOs (May 2013-April 2014)

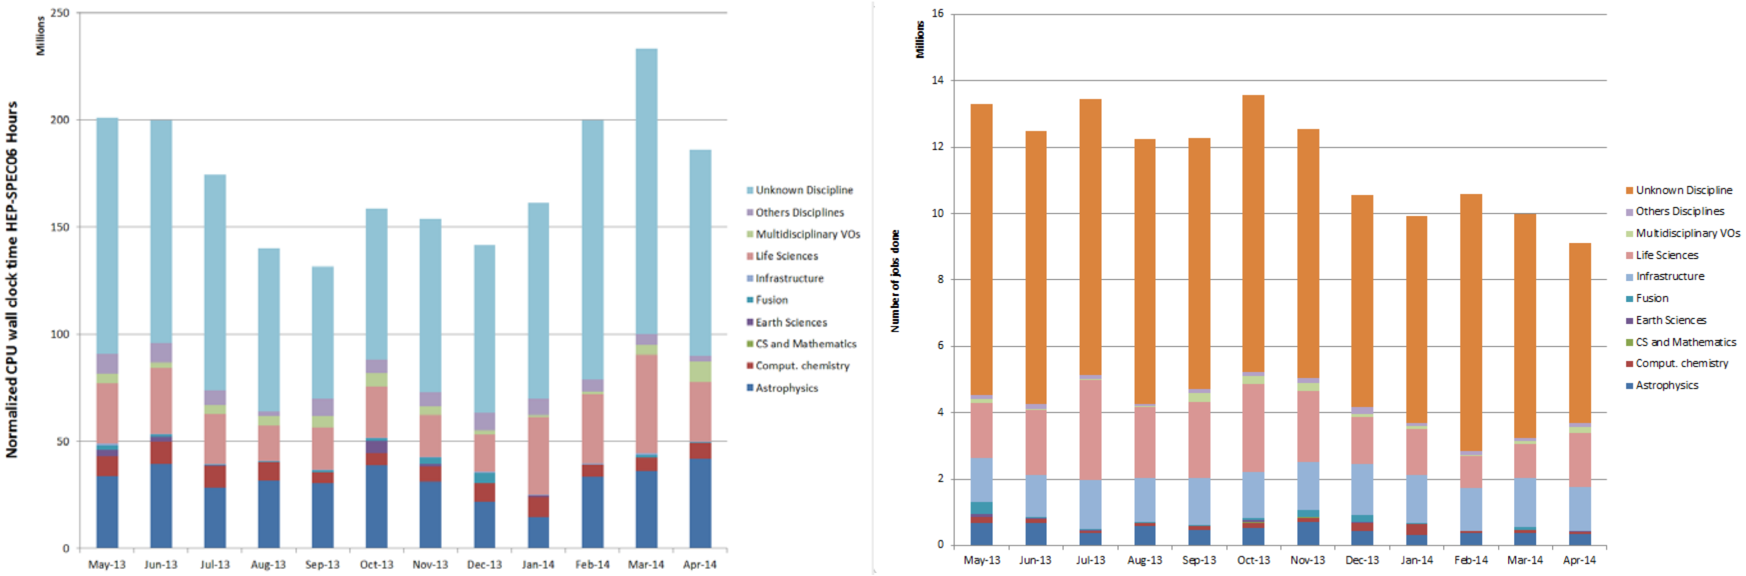

Figure 12. Used normalized CPU wall clock time (left) and number of jobs done (right) across disciplines during PY4.

5 RESOURCE USAGE

EGI accounting information is gathered and stored centrally for display through the accounting portal[10]. Accounting information is aggregated by Operations Centre, whose list is obtained from GOCDB.

Table 10. Annual compute resource usage (yearly figures)

	PY3	PY4 Value/Target
Total normalized CPU wall clock time consumed (Billion HEP-SPEC 06 hours)	12.01	15.1
Jobs per year (Million)	507.2	587.8
Average number of Jobs per day (Million)	1.43	1.61/1.6

The overall quantity of computing resources used in PY4 amounts to 15.1 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY3 amounted to 12.01 Billion HEP-SPEC 06 Hours) as shown in Table 10. The PY4 workload was generated by 587.8 Million jobs, which amounts to an average of 1.61 Million job/day.

PY4 usage expressed in HEP-SPEC 06 Hours of CPU wall time across the various resource infrastructures of EGI is plotted in Figure 13, 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. The top infrastructures have not changed since PY3. Usage distribution naturally reflects availability of installed capacity (Section 3), however the level of multidisciplinary support varies considerably across the infrastructures. Figure 144 plots the distribution of used HEP-SPEC 06 CPU wall clock hours of non-HEP user communities. NGI_FRANCE is the infrastructure with the largest absolute amount of resources used by non-HEP communities with more than 267 Million CPU wall time hours, followed by NGI_IT, NGI_DE, NGI_UK and NGI_NL.

The Figure 15 shows how support of HEP is dominant in large resource infrastructures, while other disciplines dominate in various countries in Eastern-South 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 PY4, is that the fraction of non-HEP usage will increase in future years.

[10] http://accounting.egi.eu/egi.php

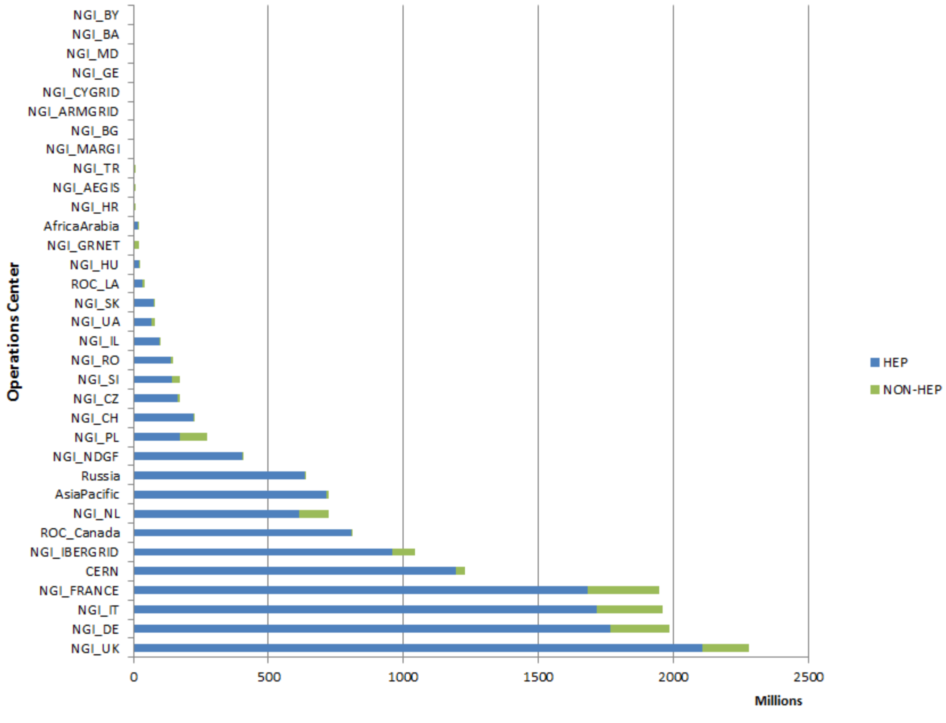

Figure 13. (HEP-SPEC 06 hours) from May 2013 to April 2014 (source: accounting portal). HEP usage is displayed in blue while the aggregated usage of non-HEP disciplines is in green

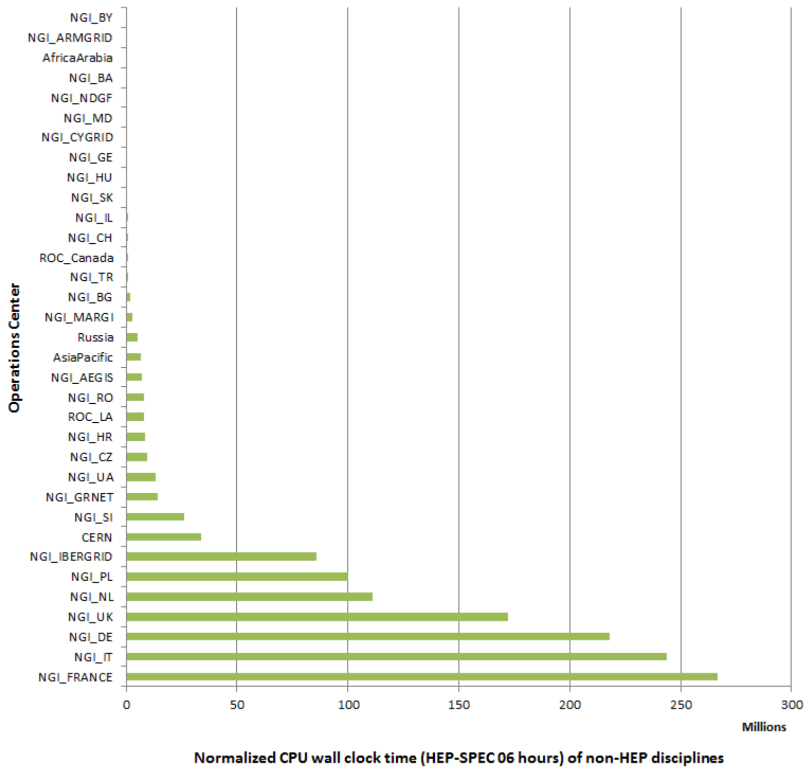

Figure 14. Distribution across EGI Operations Centres of aggregated usage of non-HEP disciplines (CPU wall clock time in HEP-SPEC 06 hours) from May 2013 to April 2014 (source: accounting portal).

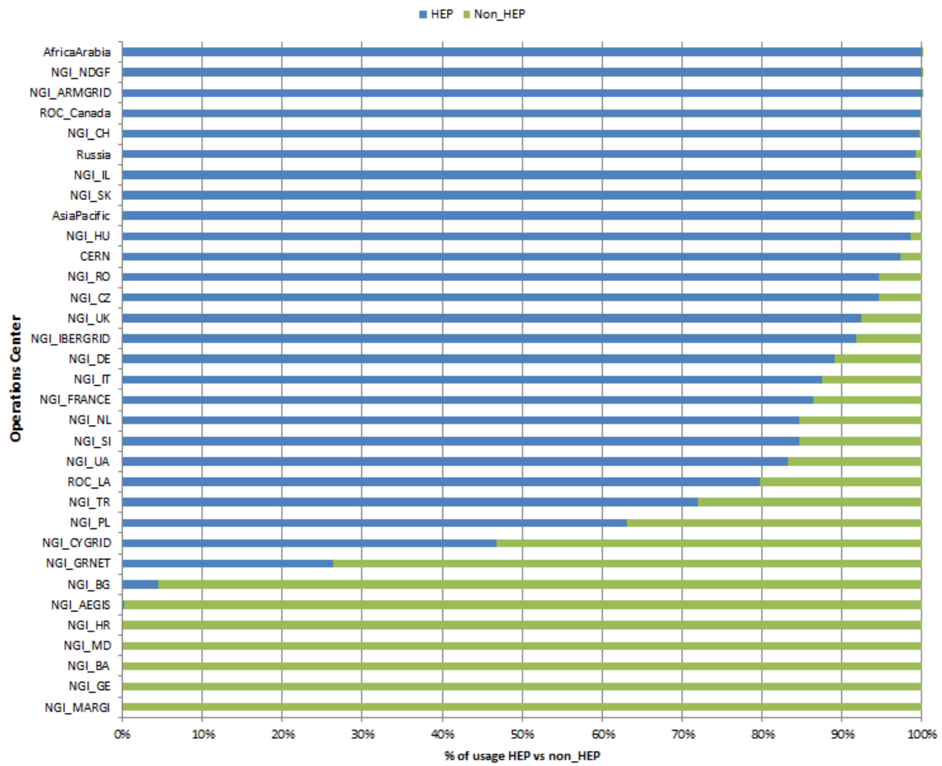

Figure 15. Distribution of resource usage (%) across HEP and non-HEP disciplines from May 2013 to April 2014 (source: accounting portal).

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 calculating 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[11]
    - Minimum Availability: 80%
    - Minimum Reliability: 85%
    - Reports: https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Resource_Centres
    - During PY4 it was agreed to increase the threshold to Availability: 80% and Reliability: 85%. Technical implementation of this change will take place during PY5.
- Resource infrastructure Providers[12]
    - Minimum top-BDII Availability: 99%
    - Minimum top-BDII Reliability: 99%
    - Maximum Regional Operator on Duty Performance Index: 10
    - Reports: https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Resource_infrastructures_Providers
- EGI.eu[13]:
    - Depending on type of service different service targets were defined:

---

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

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

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

Table 11 EGI.eu Service Level Targets

Service	Targets		
	Other	**Availability/Reliability** (as a percentage per month)	**Support priority**
1st Level Support			Maximum response 1 working hour
2nd Level Support			Advanced
Accounting Portal		Minimum availability: 99% Minimum reliability: 99%	Medium
Metrics Portal		Minimum availability: 99% Minimum reliability: 99%	Best effort
Accounting Repository		Minimum availability: 99% Minimum reliability: 99%	Medium
Collaboration Tools		Minimum availability • DNS: 99% • Other: 90% Minimum reliability: 99%	Medium
Incident Management Helpdesk		Minimum availability: 99% Minimum reliability: 99%	Medium
Monitoring Central Services		Minimum availability: 99% Minimum reliability: 99%	Medium
Operations Portal		Minimum availability: 99% Minimum reliability: 99%	Medium
Service Registry GOCDB		Minimum availability: 99% Minimum reliability: 99%	Medium
Software Provisioning Infrastructure		Minimum availability • UMD repositories, web front-end, the community repository: 90% • The other components: 75% Minimum reliability: 90%	Medium
Acceptance Criteria	Incremental definition • New document version is produced every year, following two public drafts. Verification of acceptance criteria • The verification		Base

	activities must support the UMD releases. The estimated number of products to verify in one year is 200 PPA.		
Catch All Services			Medium
Message Broker Network		Minimum availability: 95% Minimum reliability: 95%	Medium
Operations Support			Medium
SAM Central Service		Minimum availability: 99% Minimum reliability: 99%	Medium
Security Monitoring And Related Support Tools		Minimum availability: 99% Minimum reliability: 99%	Medium
Stage Rollout			Base
Security Coordination			Medium

- EGI.eu service level reports are currently under development and are expected to be available in PY5.
- The support priority is described under https://wiki.egi.eu/wiki/FAQ_GGUS-QoS-Levels#Levels_of_service

6.2 RC Performance

6.2.1 Availability and Reliability

Table 12. EGI-wide Availability and Reliability and the related project metric target.

EGI Average Monthly Reliability	May 2013-January 2014	Y4 Target
Reliability	96.25 %	97,5%
Availability	96.60 %	97

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. In January 2014 it was agreed to increase the threshold to Availability: 80% and Reliability: 85%. Technical implementation of this change will take place during PY5.

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), during May 2012 – February 2013 it slightly decrease to 93.74%/94.86%. This was probably related to the maintenance work at RCs to upgrade their software infrastructure. In the period March 2013 – January 2014 it increased to 96.25%/96.6%. The PY4 Reliability target (97.5%) was not met with a small deviation (-1,3%) – see Table 12.

The trend of the overall EGI RC availability and reliability is shown in Figure 16, 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.

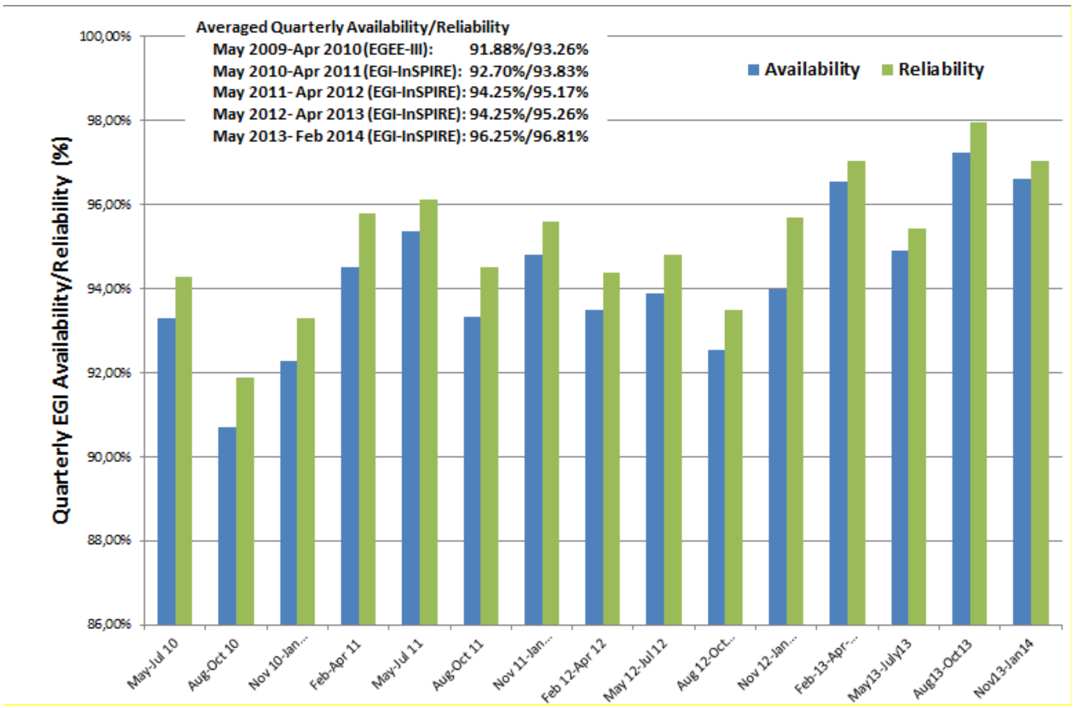

Figure 16. Quarterly availability and reliability of resource centres averaged across EGI from May 2010 to end of PQ15. Source: Availability and reliability monthly reports.

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

6.3 RP Performance

Table 13. Yearly average availability and reliability of NGI core middleware services and core operations tools services (May 2013-January 2014)

NGI core middleware services	May 2013-January 2014	PY4 Target
Reliability (MSA1.Operations.4)	99.3 %	99.8%
Availability (MSA1.Operations.4)	98.7 %	99.6%
NGI core operations tools		
Reliability (MSA1.Operations.3)	98.68%	-
Availability (MSA1.Operations.3)	98.21%	-

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.

In PY4 data concerning performance of NGI core operations tools have started to be collect on monthly basis. NGI core operations tools are those ones on which infrastructure reporting and monitoring rely on and due to that fact they are essential to be able to track performance of the infrastructure.

The average reliability performed by NGI core middleware services target was not met with a small deviation (0.5%) as shown in Table 13.

6.3.1 Availability and Reliability

Current availability and reliability reports for NGI core middleware services 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[14].

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

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

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

As shown in Figure 17, the performance of NGI core middleware 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. In PY4 the performance is stable.

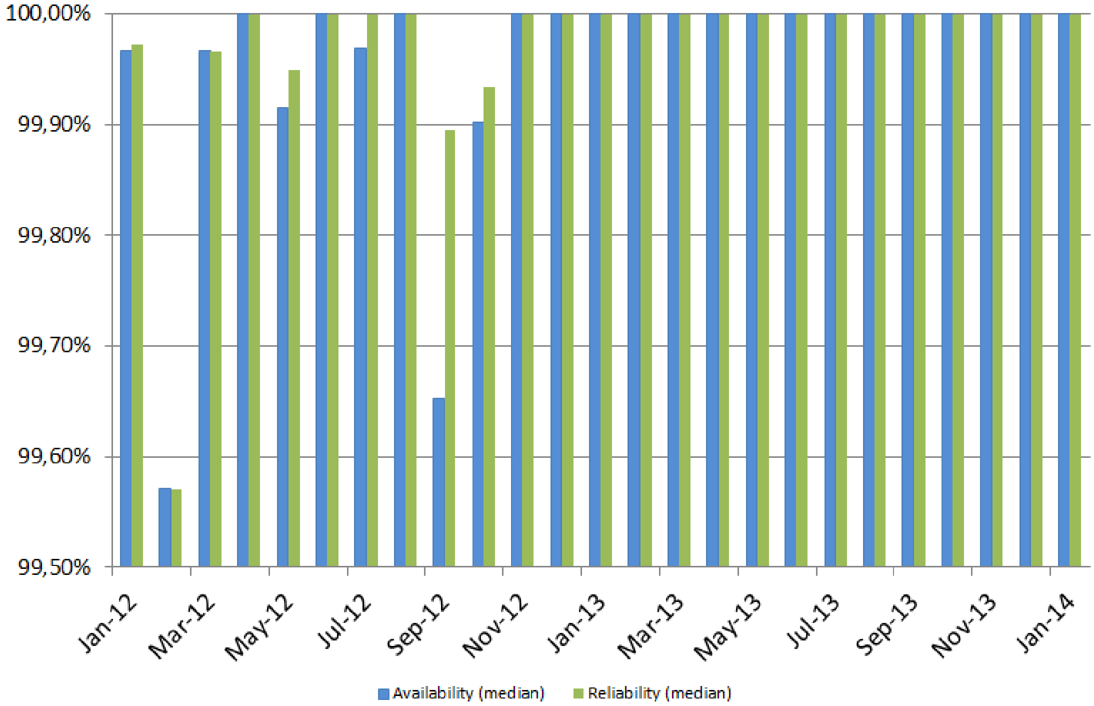

Figure 17. Median of NGI core middleware monthly Availability and Reliability performance (top-BDII service) – Jan 2012 – Jan 2014.

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 use the EGI Catch All top-BDII server provided by Greek JRU[15].

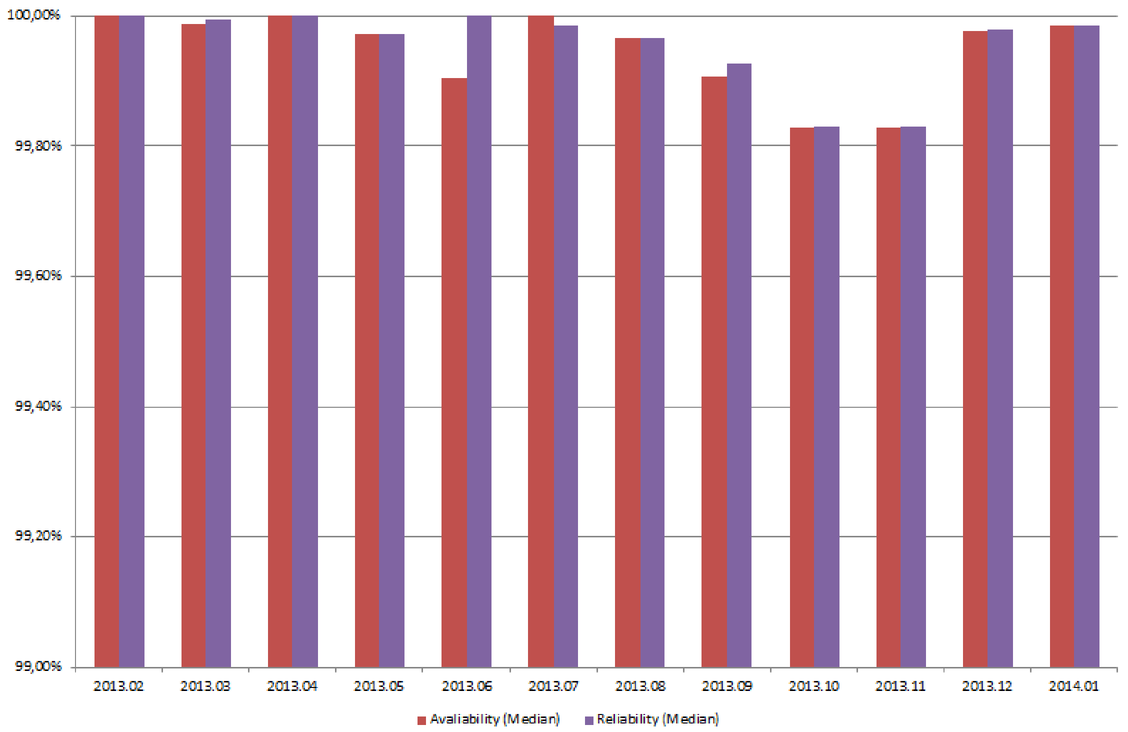

Figure 18 Median of NGI core operations tools monthly Availability and Reliability performance (SAM service) – Feb 2013 – Jan 2014.

6.4 EGI.eu Performance

Table 14. Yearly average availability and reliability of EGI.eu critical and non-critical central operations tools (PQ15)

EGI.eu critical central operations tools Monthly Performance (MSA1.Operations.6a)	**May 2013-March 2014**	**PY4 Target**
Reliability	98.01%	99.8%
Availability	97.95%	-
EGI.eu non-critical central operations tools Monthly Performance (MSA1.Operations.6b)		
Reliability	99.47%	-
Availability	99.47%	-

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

Monitoring of EGI.eu Core Infrastructure Platform was rolled to production in November 2012. In order to do so, a new central SAM instance was rolled to production to monitor these tools and various user community services.

Services have been divided into two categories:

- Critical central operations tools:
    - Monitoring Central Services
    - Accounting Repository
    - Incident Management Helpdesk
    - Operations Portal
    - Service Registry GOCDB
    - Message Broker Network
- Non-critical central operations tools
    - Metrics Portal
    - Accounting Portal
    - Application database
    - CRM

Critical central operations tools have been detected as essential for providing reliable infrastructure.

Availability statistics of these tools are now accessible through the MyEGI portal[16].

Based on new EGI.eu OLA framework the list of monitored operations tools which will be included into EGI.eu performance metrics will be extended in PY5.

The average availability and reliability performed by these EGI.eu tools in PQ15 is indicated in Table 14 and not met with a small deviation the PY4 target (99.8%). The performance was affected by some instability experienced by the Message Broker Network. The performance of this tool is being monitored on a monthly basis. From PY5 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.

6.5 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 19), 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. To consider that when the infrastructure is considered unavailable in this algorithm, there are most probably still little less than the 80% of the resources available for the users to run their activities.

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

Figure 19 shows the monthly Availability for April 2014 of the 10 most active VOs (atlas, cms, alice, lhcb, biomed, auger, ilc, vo.cta.in2p3.fr, compchem and theophys). The median of the monthly Availability for all VOs is 98.54%, while the average is 97.45%.

This prototype will be validated for all months of PY4 and will be available in production in PY5 after a testing phase as a module of the Operations Portal.

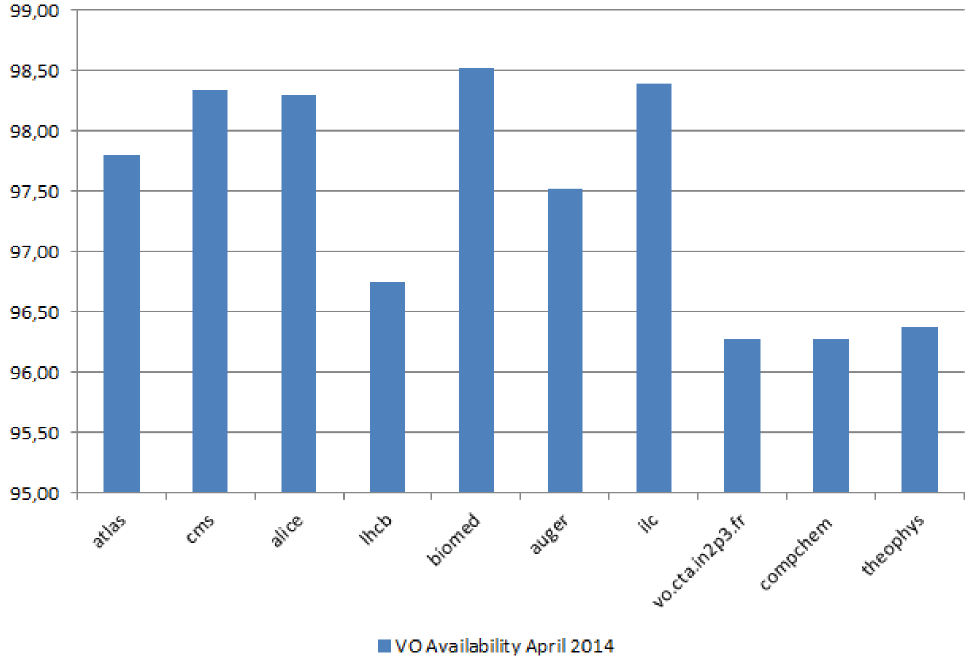

Figure 19. Availability of the top-10 active VOs (April 2014)

7 GRID SERVICES

In this section we review the status of deployment of different software platforms across EGI. As indicated in Table 15, 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[17], and Desktop Grid software released and supported by the EDGI project[18]. 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[19].

7.1 Integrated Software Platforms

Table 15. Deployment of integrated software platforms across EGI

Integrated Grid Platform	Number of countries	Countries
ARC	13	Denmark, Estonia, Finland, Germany, Latvia, Lithuania, Norway, Slovakia, Slovenia, Sweden, Switzerland, Ukraine, UK
Desktop Grid	1	Hungary
GLOBUS	5	GridFtp: Germany, UK, Poland, Belarus GRAM: Germany, Netherlands
QosCosGrid (QCG)	2	Poland, Belarus
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 services originating from the gLite distribution (now unsupported) and EMI and now supported and distributed through UMD releases are deployed by the majority of the production RCs. The number of operations centres supporting other stacks slightly increased during PY4: as shown in Figure 20, eight 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, NGI_SK, NGI_UA and NGI_UK.

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

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

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

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

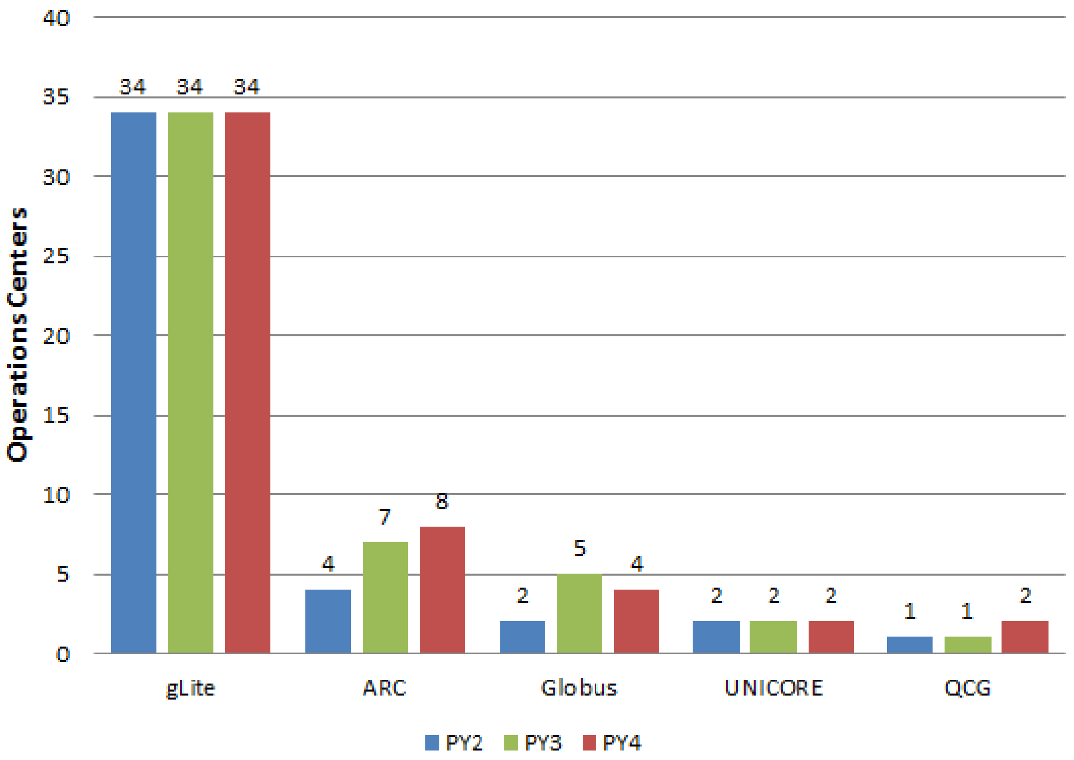

Figure 20. Deployment of the five reference grid middleware stacks across the EGI-InSPIRE operations centres, March 2013 and March 2014(source GOCDB).

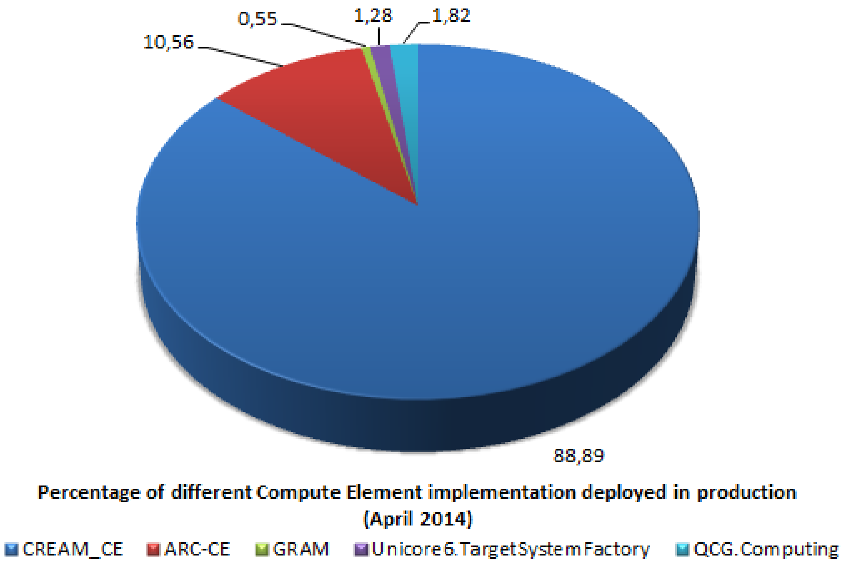

Figure 21. Number of instances of the different implementations of the compute capability, across the EGI_InSPIRE partners and the integrated Resource infrastructure Providers, April 2014 (source: GOCDB)

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 21 shows this distribution: CREAM-CE is in production in the 88.89% of the infrastructure, ARC-CE is second in deployment (10.56%) followed by QCG.Computing (1.82%), Unicore6.TargetSystemFactory (1.28%) and GRAM (0.55).

7.2 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 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 22 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 2014 the EGI integrated infrastructure comprises 313 core services: 57 VOMS instances, 126 WMS, 30 LFC and 100 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.

(a) Overall distribution

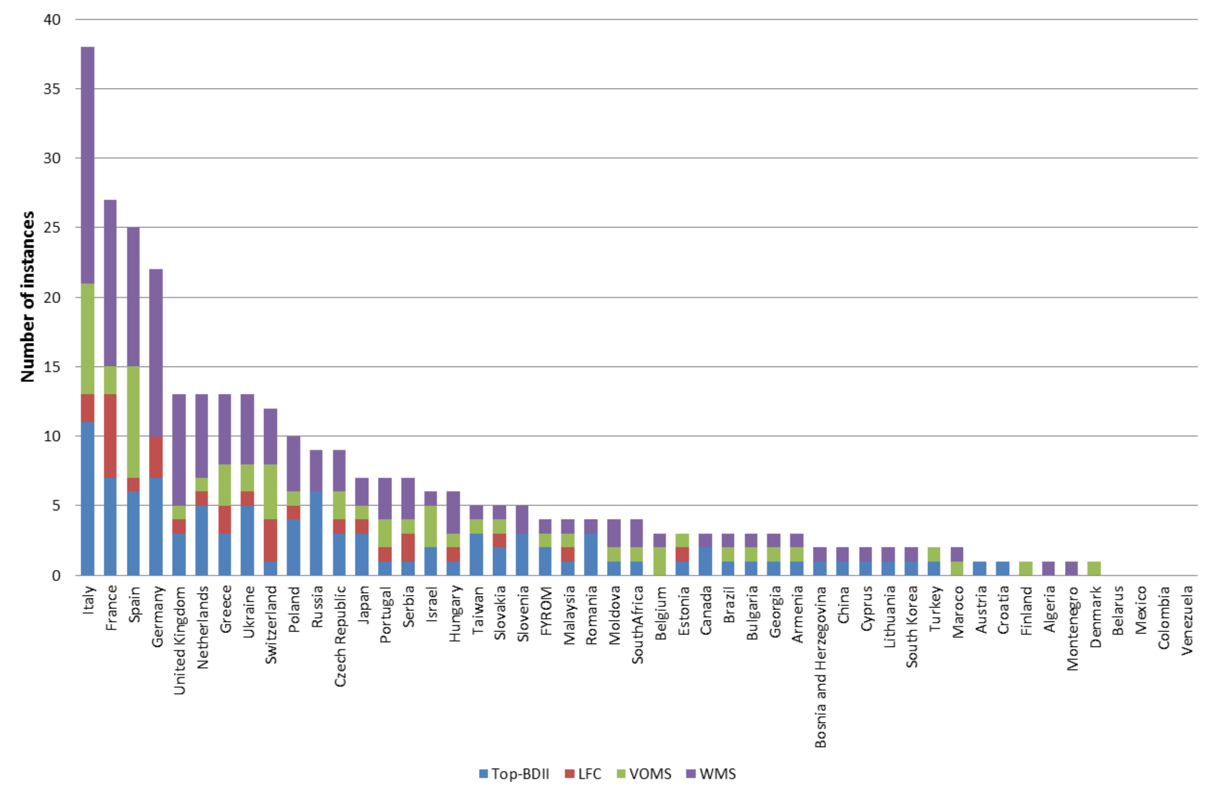

(b) Workload Management (WMS)

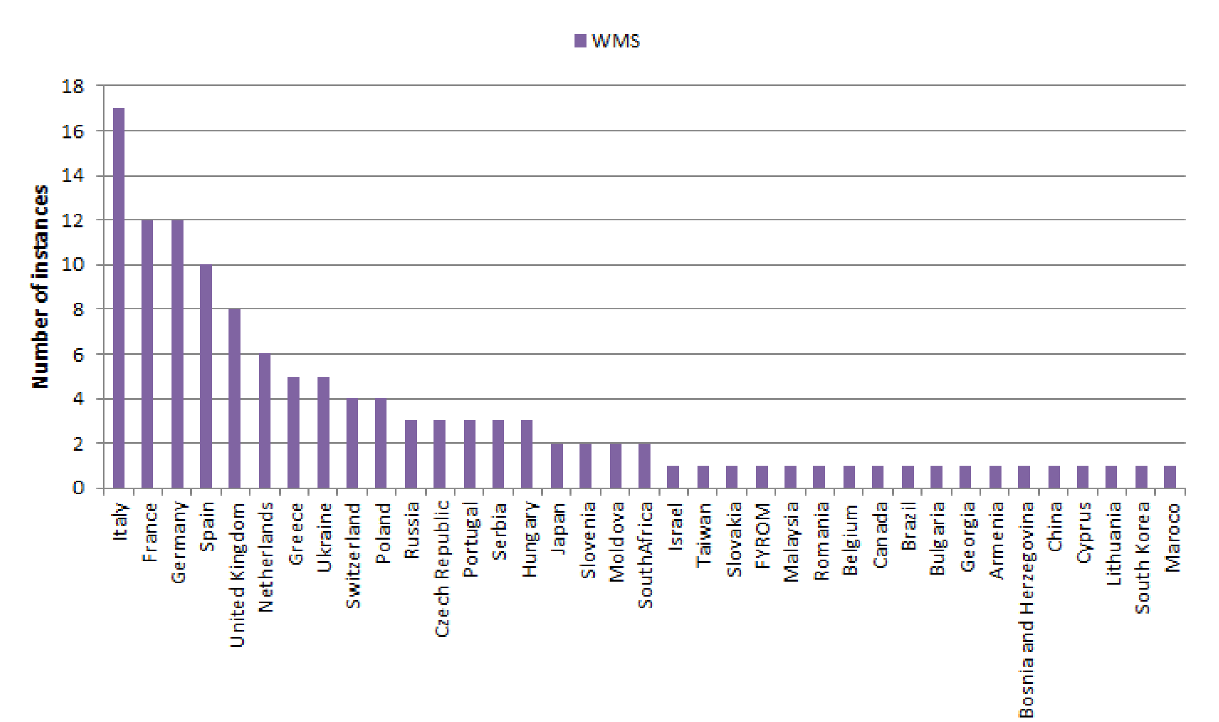

(c) VO membership (VOMS)	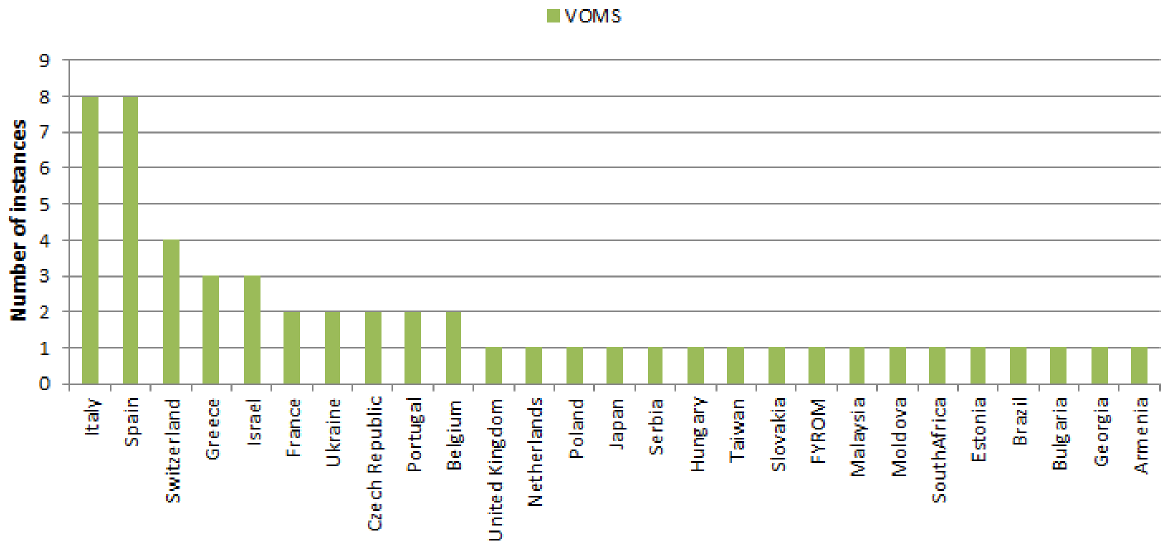
(d) file catalogue (LFC)	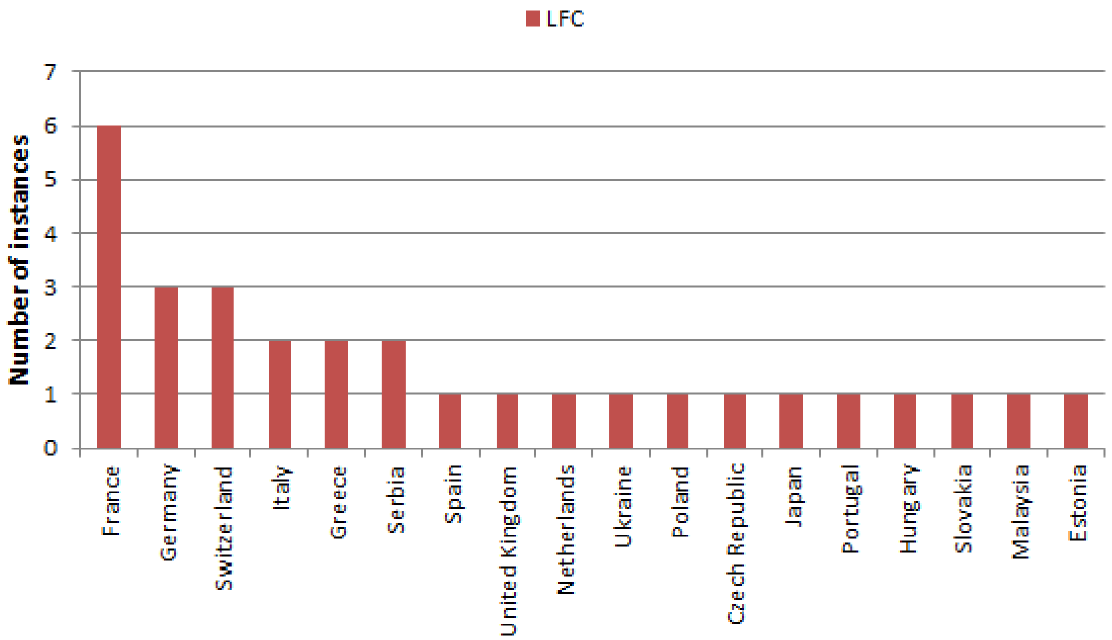

(e) Information system (Top-BDII)	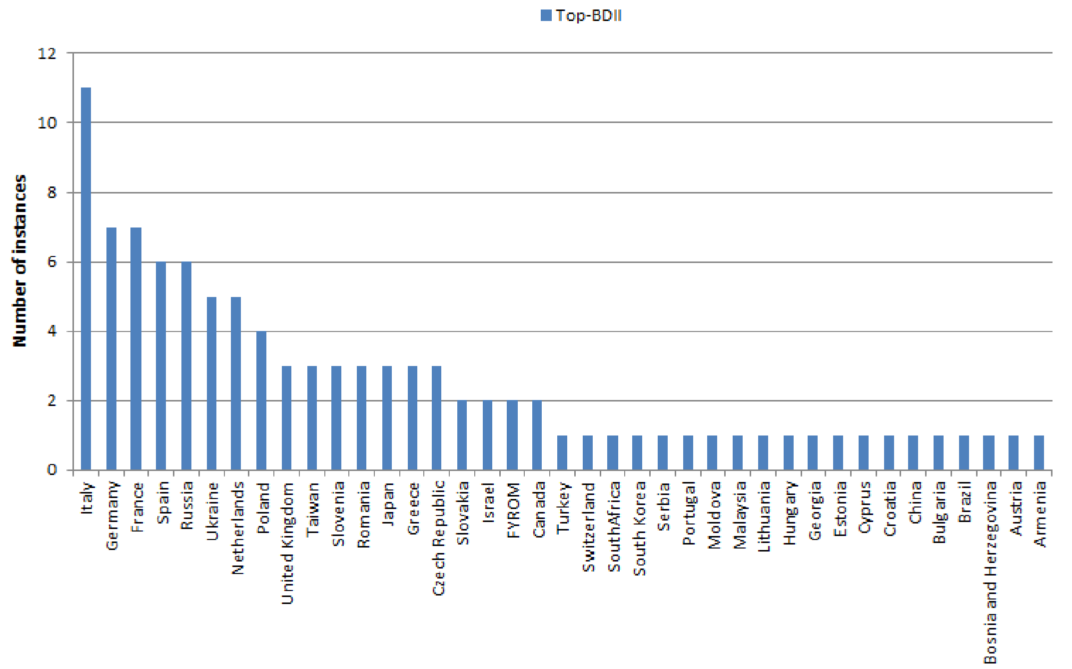

Figure 22, Number of core services instances deployed within the EGI-InSPIRE integrated infrastructure. Source: GOC DB (March 2014).

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 16: Resource providers participating to the Federated Clouds in being in production infrastructure (April 2014)

Resource centre name	Number of cores available	Amount of disk space available	Cloud middleware deployed
100 Percent IT Ltd	120	16TB	Open Stack
UNIZAR/BIFI	360	500 GB	Open Stack
Cyfronet	200	20 TB	Open Stack
FCTSG	296	500 GB	Open Nebula
CESNET	240	960 GB	Open Nebula
GWDG	192	40 TB	Open Nebula
CSIC/IFCA	144		Open Stack
GRNET	20	400 GB	Synnefo
II SAS	96	16 TB	Open Stack
INFN-Catania	120	6 TB	Open Stack, Open Nebula
KTH	192	6 TB	Open Nebula
INFN-Bari	300	50 TB	Open Stack
MTA SZTAKI	128	6 TB	Open Nebula

Table 17: Resource providers under the process of integration in the Federated Clouds test-bed (April 2014)

	Number of cores potentially contributed to the testbed	**Amount of disk space potentially contributed to the test-bed**	**Cloud middleware deployed**
CRNS/IN2P3-CC			Open Stack
TUBITAK ULAKBIM	112	6 TB	Open Stack
UKIM	216	17 TB	Open Nebula
KISTI	64	6 TB	Open Stack

The EGI Federated Cloud testbed has been integrated into the production infrastructure in PY4.

Table 16 contains the resource providers participating to the activities of the Federated Cloud task force and already integrated in production infrastructure. As shown in the list the most common cloud middleware solutions are Open Nebula and Open Stack. All Resource Providers are being monitored with a cloud 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 17 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. 19 more Resource Centres are interested to join Federated Cloud activity.

The two table reports a summary of the resources contributed to the Federated Cloud 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.

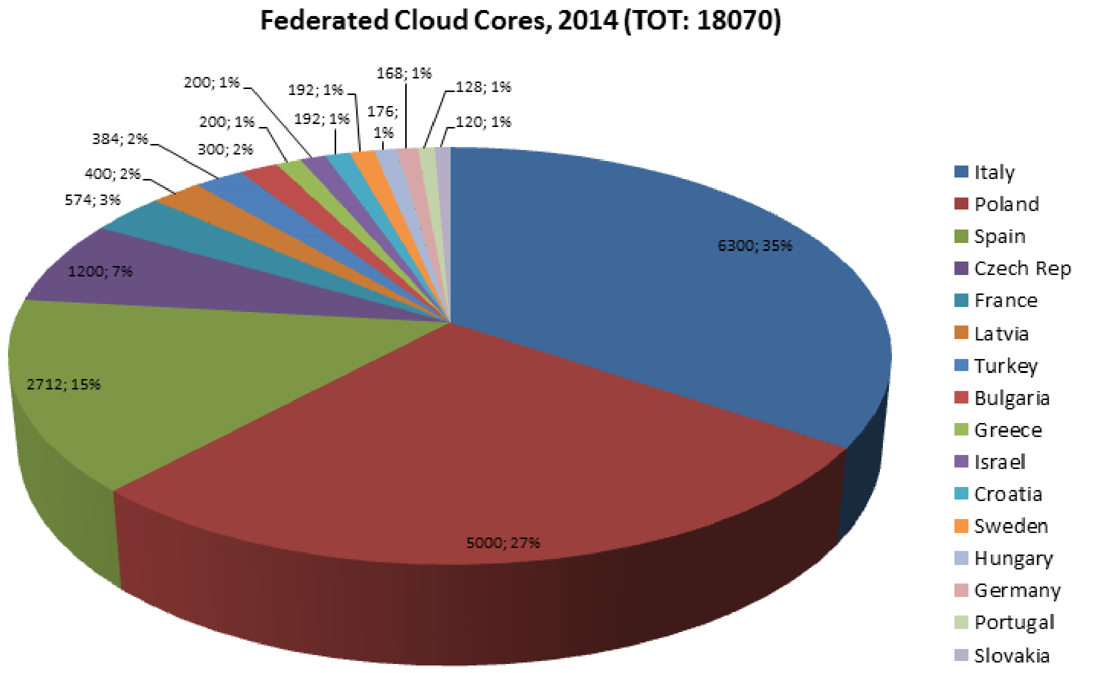

Figure 23 Federated Cloud Cores distribution per country (April 2014)

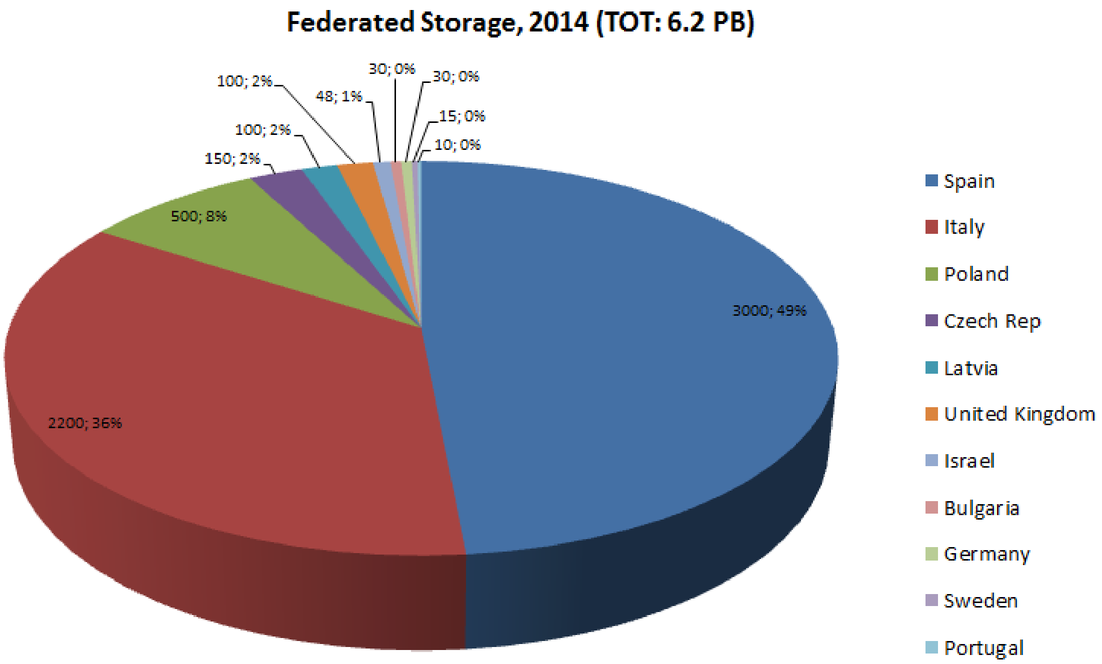

Figure 24 Federated Cloud Storage distribution per country (April 2014)

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 18. Overview of EGI-InSPIRE Staged Rollout metrics.

Metric	PQ13	PQ14	PQ15	PQ16
Number of staged rollouts	53	33	38	36
Number of components	39	24	19	22
Number of sites	20	17	16	12

As shown in Table 18 the largest number of products was tested in PQ13, as EMI released a major release and a very big update before the conclusion of the project. This number was gradually reduced. The number of participating EAs currently amount to 80 teams.

The number of tests performed from PQ13 to PQ16 by NGIs and EIROs is plotted in Figure 25.

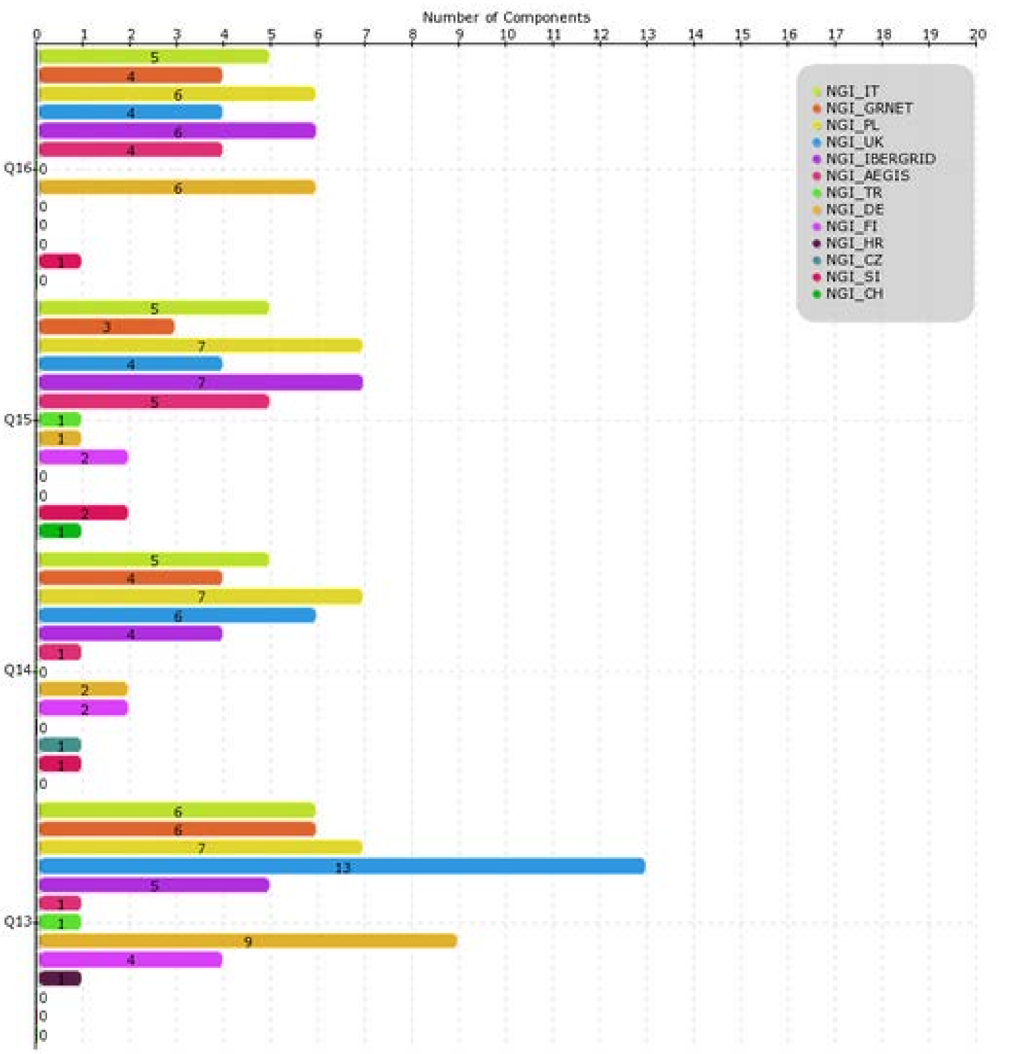

Figure 25. Number of Staged Rollout tests performed from PQ13 to PQ16 by NGIs/EIROs. (source: Staged Rollout portal)

10 SOFTWARE SUPPORT

Software Support was performing the following essential functions throughout PY4:

- 1st level support, i.e., Ticket triage and assignment for dispatching of tickets to the appropriate Support Units within GGUS
- 2nd level software support, encompassing both grid middleware and operational tools
- Ticket oversight and follow up

The volumes of tickets processed at different levels were as follows:

Metric	PQ13	PQ14	PQ15	PQ16
Assigned to DMSU	125	100	122	86
Forwarded to 3rd level	95	62	87	66
Solved by DMSU	30	34	29	22

A total of 433 tickets arrived to the DMSU during PY4. On average, 27 % of the incoming tickets were resolved within DMSU. The rest were forwarded to other SUs for resolution. In Q16 software support received 86 tickets, which is a lower number compared to previous periods, indicating that the set of deployed components is in more stable state at the end of the project. Out of those, 22 were solved at the 2nd level, which is around the usual ratio.

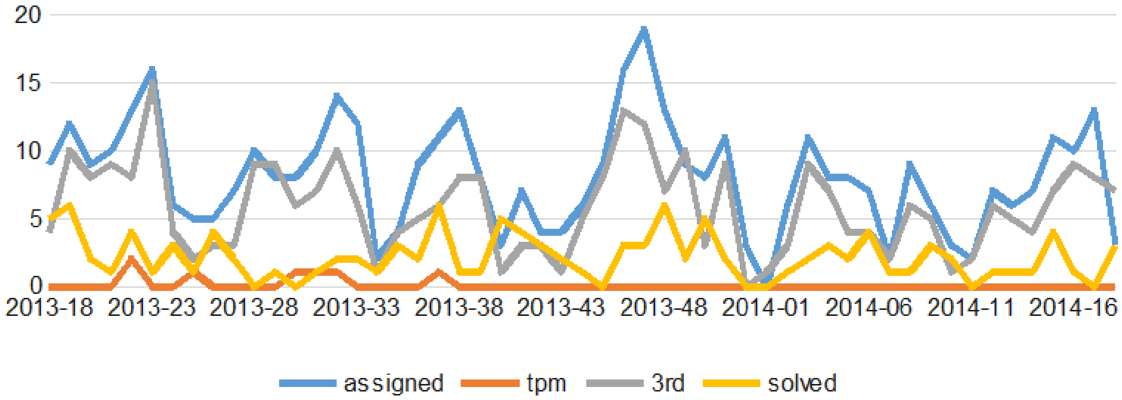

Figure 26. Weekly distribution of tickets handled by the 2nd level Support Unit

The software support unit interacts regularly with EGI Operations, identifying issues with potential broader impact on the infrastructure and describing them in a dedicated wiki page. The software support unit's 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. Software Support unit representatives also attend TCB meetings.

During PY4, EGI has called for bids to provide software support from the end of PY4 onwards. Two bidding consortia were selected. During the final months of PY4, the software support unit was preparing to hand over support responsibilities to the new partnership, ensuring that support is provided along similar lines after PY4 ends. New procedures were agreed: the 1st level will be provided in weekly shifts alternating between CESNET and IBERGRID. At CESNET the shifts are integrated into the work of its general 24/7 helpdesk.

As part of the handover effort, the list of supported software components at the 2nd level was revised and components were assigned to CESNET, IBERGRID, and other unfunded contributing partners (in case of the community platform). A few "orphaned" products were identified, and the partners have pledged to support them on a best-effort basis, bordering on regular support.

The new work model for 1st level was deployed in advance in April, no major problems appeared. Handover of 2nd level took place seamlessly at the turn of April and May

11 GLOBAL OPERATIONS TASKS REVIEW

In PY3 at the Evolving EGI Workshop held in January 2013 in Amsterdam supporting activities and related costs were discussed in depth and the impact of not doing the service was rated between critical, degradation, no growth and no impact. Part of EGI's sustainability efforts and short-to-medium term strategy is based on what the community feels is the most critical and the costs to ensure those services are carried out in the future. In PY4 the EGI Council was further detailing a plan and strategy that takes into account this costing work.

All Global Operations Tasks have being evaluated within the EGI Council for long-term strategic impact and decisions and potential sustainability models for what tasks should be offered and funding models to support them. The "Global Tasks' provided by EGI.eu partners were re-bid in 2013[20]. A transition plan has been defined in September 2013 and implemented in next six months of the project.

Last year's report (MS123), allowed understanding the cost of the various activities and services and analysing the individual areas such as operations, maintenance, and development. As the EGI Global Tasks have evolved into a more structured service portfolio, EGI.eu management was able to look at also the effort required for coordination and support of these services.

Following table provides matching between Global Operations Tasks used in PY3 Global Operations Tasks Review (MS123) and final list of Core EGI Activities, essential for providing reliable infrastructure.

Task name	Core EGI Activity	Description
Project and Programme Management	Collaborations tools	Collaborations tools part of the task has been assessed as Core Activity. Collaborations tools are services needed by the EGI back-office and supporting EGI collaboration.
Operations Coordination	Operations support	Grid oversight part of the task has been assessed as Core Activity. Operations support is auxiliary service needed by the Core Infrastructure Platform and by various operational activities of EGI. Auxiliary activities are needed for the good running of Infrastructure Services. Examples of such are activities for service level management, service level reporting, service management in general and central technical.
Security Coordination	Security Coordination	Central coordination of the security activities ensures that policies, operational security, and maintenance are

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

		compatible amongst all partners, improving integrity and availability and lowering access barriers for use of the infrastructure.
Helpdesk Support	1st and 2nd level support	First level support is responsible for ticket triage and assignment. This activity is also responsible for the coordination with teams responsible for 2nd level and 3rd level support. Software-related tickets that reach the second level of support are analysed and if necessary are forwarded to 3rd line support units only when there are clear indications of a defect (in software, documentation, etc.). 3rd level support is provided by Product teams based on agreement with EGI.eu.
Repository of Validated Software	Acceptance Criteria, Staged Rollout, Software provisioning infrastructure	Acceptance Criteria, Software Provisioning infrastructure and Staged Rollout part of the task have been assessed as Core Activity. The Acceptance Criteria are the functional and non-functional requirements that a product must fulfil to be released in UMD, these include generic requirement applicable to every product, and specific requirements applicable to the capabilities supported by a component. The Software Provisioning infrastructure provides the technical tools to support the UMD release process from pulling packages from the developers' repositories to the build of a release. The Staged Rollout is an activity by which certified updates of the supported middleware are first tested by Early Adopter (EA) sites before being made available to all sites through the production repositories. This procedure permits to test an update in a production environment that exposes the product to more heterogeneous use cases than the certification and verification phase. This allows the discovery of potential issues and potentially to add mitigation information to the UMD release notes.
Accounting Portal and	Accounting Repository,	Task has been divided into Accounting Repository and

Repository	Accounting and Metric Portal	Accounting and Metric Portal. The Accounting Repository stores user accounting records from various services offered by EGI. It is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI. The EGI Accounting Infrastructure is distributed. At a central level it includes the repositories for the persistent storage of usage records. The central databases are populated through individual usage records published by the Resource Centres, or through the publication of summarised usage records. The Accounting Infrastructure is essential in a service-oriented business model to record usage information. The Accounting Portal provides data accounting views for users, VO Managers, NGI operations and the general public. The Accounting Portal is part of the EGI Core Infrastructure Platform which supports the daily operations of EGI. The EGI Accounting Infrastructure is distributed. At a central level it includes the repositories for the persistent storage of usage records. The central databases are populated through individual usage records published by the Resource Centres, or through the publication of summarised usage records. The Accounting Infrastructure is essential in a service-oriented business model to record usage information.
Catch-all Grid Services for small user communities	Catch-all services	Catch-All services are auxiliary services needed by the Core Infrastructure Platform and by various operational activities of EGI. Auxiliary services and activities are needed for the good running of Infrastructure Services. Examples of such services are VOMS service and VO membership management for infrastructural VOs (DTEAM, OPS), the provisioning of middleware services needed by the monitoring infrastructure (e.g. top-BDII and WMS), and catch-all services for emerging user communities.
Development of operations monitoring probes	Monitoring Central services	Task has been merged with Monitoring Central services activity.
GOCDB	Service Registry (GOCDB)	Service Registry (GOCDB) is a central registry to record information about different entities such as the Operations Centres, the Resource Centres, service endpoints and the contact information and roles of people responsible for operations at different levels. GOCDB is a source of

		information for many other operational tools, such as the broadcast tool, the Aggregated Topology Provider, the Accounting Portal, etc. GOCDB is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI.
Incident Management Tool (EGI Helpdesk)	Incident Management (Helpdesk)	Incident Management (Helpdesk) is the central helpdesk provides a single interface for support. The central system is interfaced to a variety of other ticketing systems at the NGI level in order to allow a bi-directional exchange of tickets. GGUS is part of the EGI Collaboration Platform and is needed to support users and infrastructure operators.
Message Broker Network	Message Broker Network	The message broker network is a fundamental part of the operations infrastructure ensuring message exchange for monitoring, the operations dashboard and accounting. As such it is a critical infrastructure component whose continuity and high availability configuration must be ensured. The Message Broker Network is part of the EGI Core Infrastructure Platform which is needed to support the running of tools used for the daily operations of EGI.
Metrics Portal	Accounting and Metric Portal	The Metrics Portal aggregates metrics from the EGI Infrastructure from activity leaders and NGI managers in order to quantify and track the infrastructure evolution. This task was not recognized as critical but included with best effort support to Accounting Portal provisioning.
Operational Tools and Meta-service Monitoring (Ops-Monitor)	Monitoring Central services	Monitoring Central Services is supporting monitoring of activities to be conducted centrally, like monitoring of e.g. UserDN publishing in accounting records, GLUE information validation, software versions of deployed middleware, security incidents and weaknesses and EGI.eu technical services. Central Monitoring Services is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI.
Operations Portal	Operations Portal	The Operations Portal provides VO management functions and other capabilities which support the daily operations of EGI. It is a central portal for the operations community that offers a bundle of different capabilities, such as the broadcast tool, VO management facilities, a security dashboard and an operations dashboard that is used to display information about failing monitoring probes and to open tickets to the Resource Centres affected. The dashboard also supports the central grid oversight activities. It is fully interfaced with the EGI Helpdesk and the monitoring system through messaging. It is a critical component as it is used by all EGI

		Operations Centres to provide support to the respective Resource Centres. The Operations Portal provides tools supporting the daily running of operations of the entire infrastructure: grid oversight, security operations, VO management, broadcast, availability reporting.
Security monitoring tools	Security monitoring and related support tools	Security monitoring and related support tools are part of the EGI Core Infrastructure Platform which supports the daily security operations of EGI. EGI is an interconnected federation where a single vulnerable place may have a huge impact on the whole infrastructure. In order to recognise the risks and to address potential vulnerabilities in a timely manner, the EGI Security Monitoring provides an oversight of the infrastructure from the security standpoint. Also, sites connected to EGI differ significantly in the level of security and detecting weaknesses exposed by the sites allows the EGI security operations to contact the sites before the issue leads to an incident. Information produced by security monitoring is also important during assessment of new risks and vulnerabilities since it enables to identify the scope and impact of a potential security incident.
Service Availability Monitoring (SAM) central service	SAM central services	The Service is part of the EGI Core Infrastructure Platform which supports the daily operations of EGI. Central systems are needed for accessing and archiving infrastructure monitoring results of the services provided at many levels (Resource Centres, NGIs and EGI.EU), for the generation of service level reports, and for the central monitoring of EGI.eu operational tools and other central monitoring needs.
Tools (Grid Services) for Resource Centre certification	Catch-all services	This task has been merged with Catch-all Grid Services for small user communities.

Service level targets, which have been set for each of Core EGI.eu Services, have been described in Section 6.1 and agreed in corresponding OLA documents.

12 CONCLUSIONS

The production Infrastructure satisfactorily met the PY4 targets of the SA1 project metrics: the number of RCs integrated, compute resources offered, and the average number of Jobs per day. The contribution of EGI-InSPIRE SA1 to the accomplishment of the applicable project objectives is described[21].

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

  This objective was successfully met by completing the integration of the ArabiaAfrica Operations Center comprising 11 production RCs. In PY5 the plan is to extend the amount of cloud resources and sites from integration new resource infrastructures providers: IHEP-China, C-DAC India, IDGF. In addition with the end of PY4 production infrastructure has been extended to provide cloud resources.

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

  In PY4 SA1 has started resource allocation activity that is supposed to simplify reaching an agreement between users and Resource Centers concerning the parameters and conditions of use of resources. The Resource Allocation process is useful for users (VO representatives, individual users) because in a multi-provider EGI environment they have a single point of contact to ask for a share on resources and also beneficial for resource providers as it allows more effectively plan the use of resources and closer communication with users.

  We are also observing constant increase of infrastructure usage. The overall quantity of computing resources *used* in PY4 amounts to 15.1 Billion HEP-SPEC 06 Hours (the corresponding amount of consumed resources consumed during PY3 amounted to 12.01 Billion HEP-SPEC 06 Hours). The PY3 workload was generated by 587.8 Million jobs, which amounts to an average of 1.61 Million job/day.

  The overall compute resource utilization during PY4 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 2013-March 2014 the rate of jobs succssfully executed incrased by +10.8%, while the total normalized CPU wall time (HEP-SEPC06) incrased by +15.98%.

  While the HEP utilization is still dominating in absolute terms but decrease by -1.46% comparing to PY3 (88.34% of the total EGI consumption in normalized CPU wall time hours), a number of other communities significantly increased their yearly CPU wall time utilization: Fusion (+108.46% yearly increase), Computational Chemistry (+48.31%), Multidisciplinary (+76.13), Infrastructure (+14.75%) and unknown disciplines (+6417.66%).

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

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

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. Collaboration was established with the EUDAT and PRACE infrastructures and user communities 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*.

  During PY4 as part of EGI.eu service catalogue, Federated operations service has been defined. Federated operations brings together the operational tools, processes and people necessary to guarantee standard operation of heterogeneous infrastructures from multiple independent providers, with a lightweight central coordination. This includes, for example, the monitoring, accounting, configuration and other services required to federate service provision for access by multiple research communities. A federated environment is the key to uniform service and enables cost-efficient operations, while allowing resource centres to retain responsibility of local operations. This service is supposed to simplifies the day-to-day operations of a federated heterogeneous infrastructure avoiding duplication of costs and providing re-usable tools. In addition all activities which are part of the service were covered by signing an Operations Level Agreement document which describes expectations towards provisioning of the activity/tools.

  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. The version released in PQ13 has been tested, verified and deployed.

  A collaboration with EUDAT has been also 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. The collaboration refocused in PY4 to understand possible integration of helpdesk/support and security solutions.

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

  In PY4 SA1 successfully accomplished integration of cloud resources into production infrastructure. Work has been done is terms of adjustment of Core EGI.eu tools, activities and documentation to support other types of resources than grid.

13 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)
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)
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)
SRW	Staged Rollout: https://wiki.egi.eu/wiki/Staged-Rollout
UMD	Unified Middleware Distribution (http://repository.egi.eu/category/umd_releases/distribution/umd_1/)

14 ANNEX I. VO DISTRIBUTION PER DISCIPLINE

14.1 Astronomy Astrophysics and Astro-particle Physics

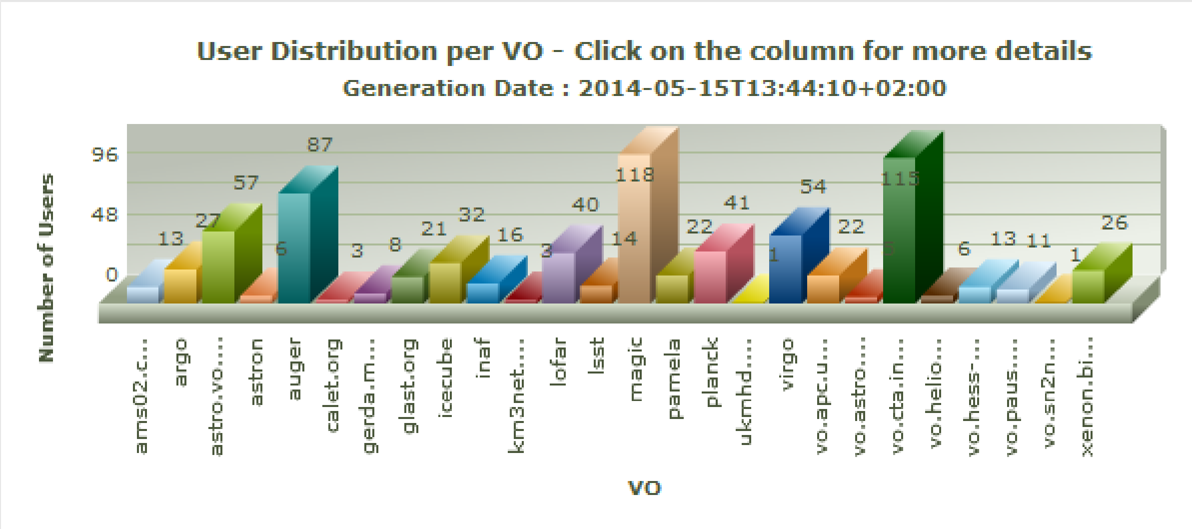

14.2 Computer Science and Mathematics

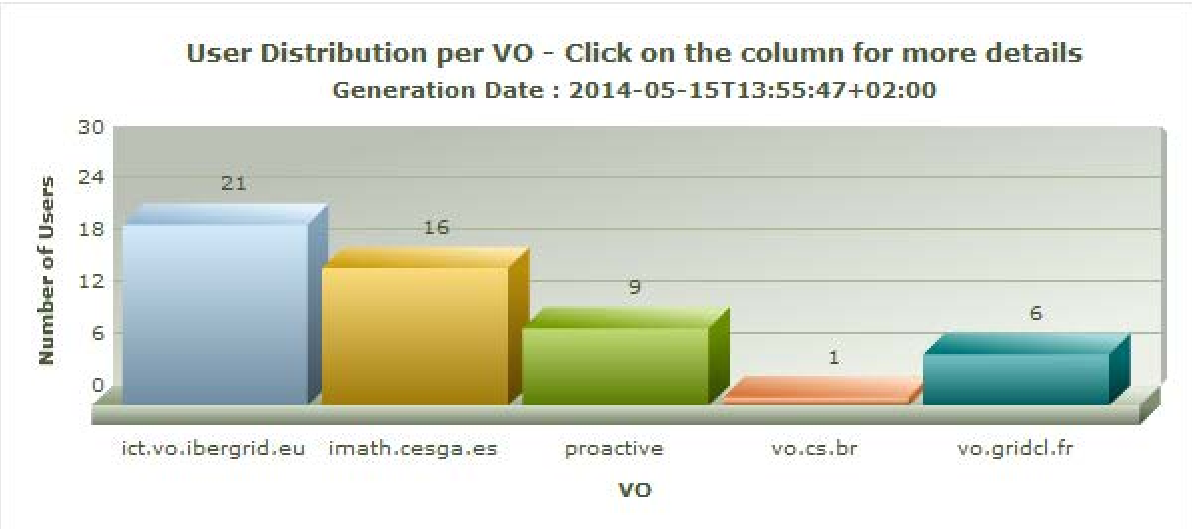

14.3 Computational Chemistry

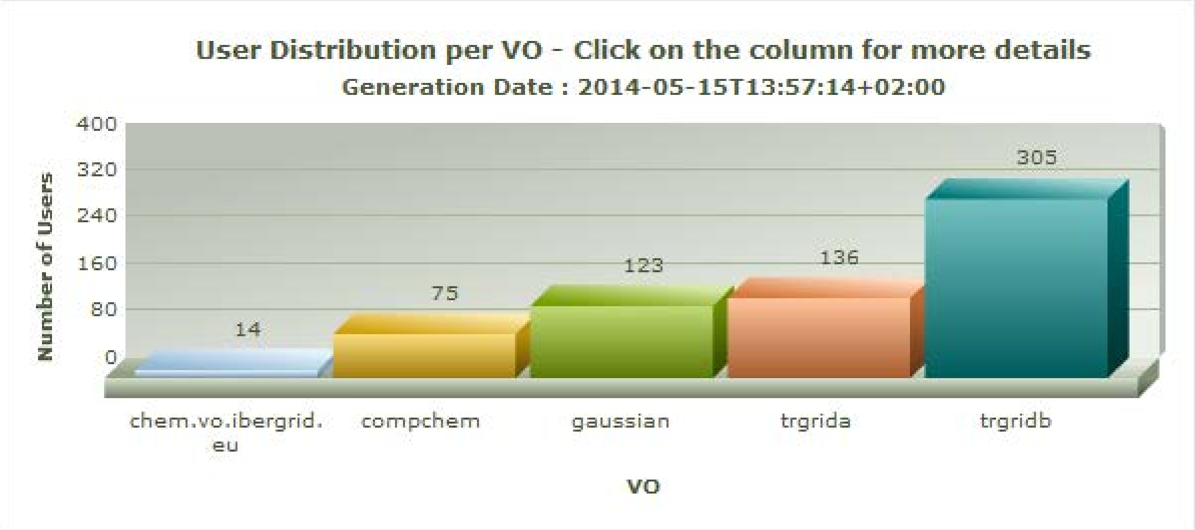

14.4 Earth Sciences

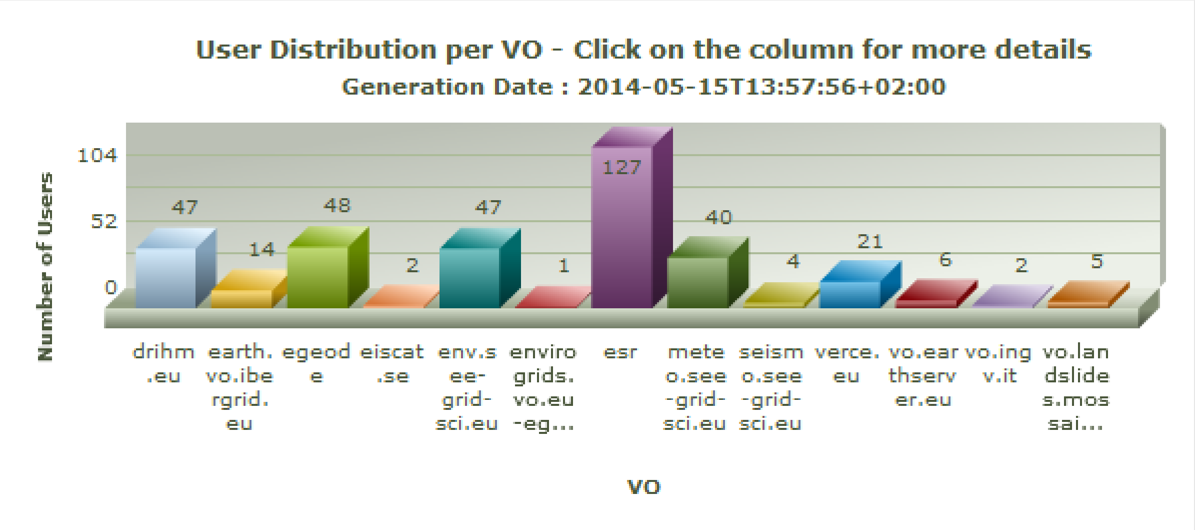

14.5 Fusion

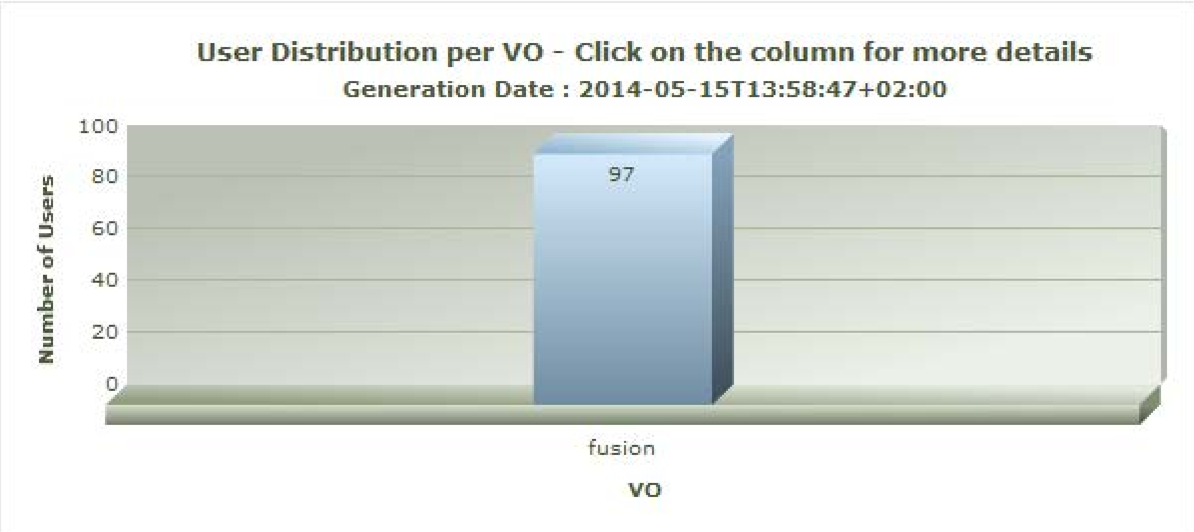

14.6 High Energy Physics

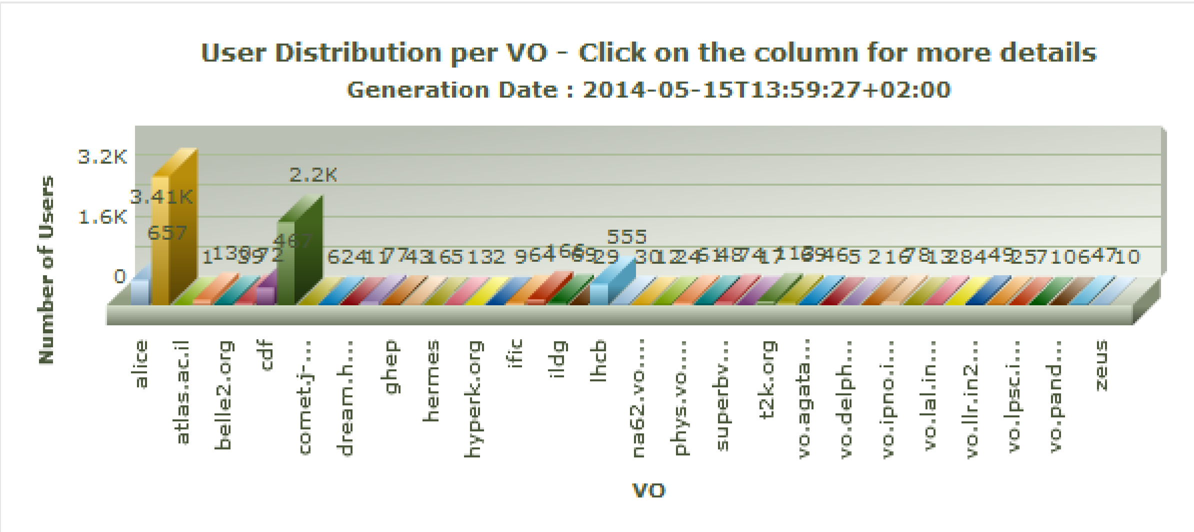

14.7 Infrastructure

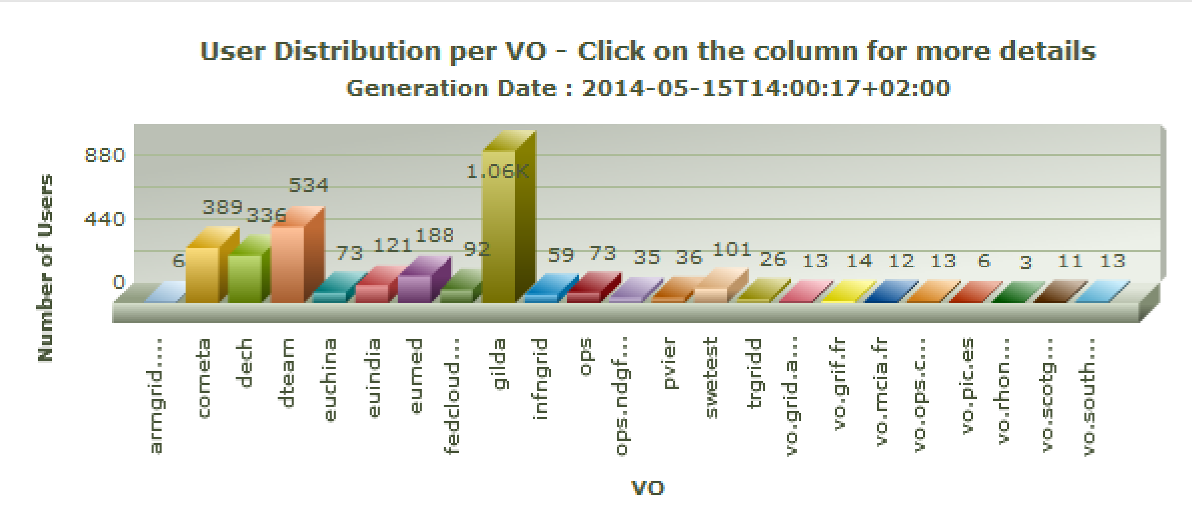

14.8 Life Sciences

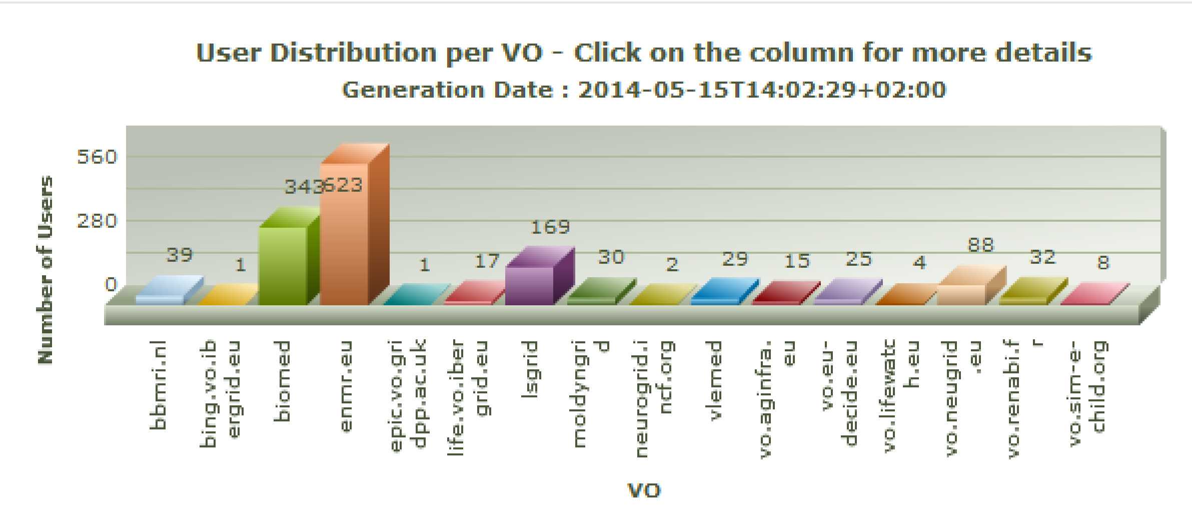

14.9 Multidisciplinary VOs

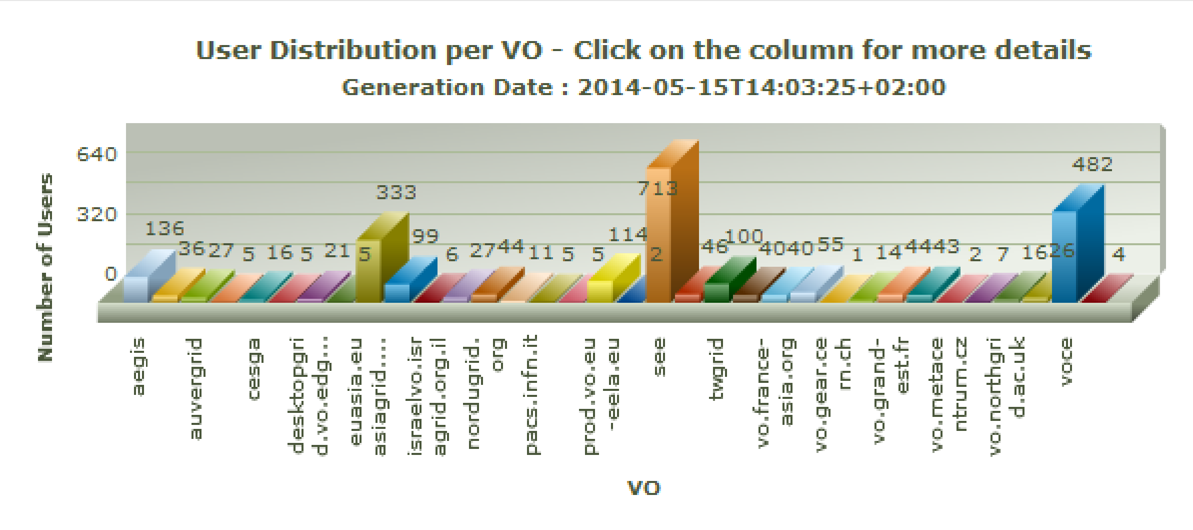

14.10 Other Disciplines