



EGI-InSPIRE

ANNUAL REPORT ON THE EGI PRODUCTION INFRASTRUCTURE EU DELIVERABLE: D4.5

Document identifier:	EGI-D4.5-v1.0
Date:	07/05/2012
Activity:	SA1
Lead Partner:	EGI.eu
Document Status:	DRAFT
Dissemination Level:	PUBLIC
Document Link:	https://documents.egi.eu/document/1059

Abstract

This document provides information on the status of the EGI Resource Infrastructure at the end of PY2. In particular, it describes the status and progress of Resource centres, Resource infrastructure Providers and Operations Centres that are contributing resources to the user community and who are responsible of the daily operations of the infrastructure. The document provides information on the amount of installed capacity provided, the status of the current EGI user base, the trends in usage, the service levels provided and the status of VO Services and grid common infrastructure services. The status of the Staged Rollout infrastructure for software testing is also presented.



I. COPYRIGHT NOTICE

Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration. EGI-InSPIRE (“European Grid Initiative: Integrated Sustainable Pan-European Infrastructure for Researchers in Europe”) is a project co-funded by the European Commission as an Integrated Infrastructure Initiative within the 7th Framework Programme. EGI-InSPIRE began in May 2010 and will run for 4 years. This work is licensed under the Creative Commons Attribution-Noncommercial 3.0 License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc/3.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, and USA. The work must be attributed by attaching the following reference to the copied elements: “Copyright © Members of the EGI-InSPIRE Collaboration, 2010-2014. See www.egi.eu for details of the EGI-InSPIRE project and the collaboration”. Using this document in a way and/or for purposes not foreseen in the license, requires the prior written permission of the copyright holders. The information contained in this document represents the views of the copyright holders as of the date such views are published.

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	T. Ferrari	EGI.eu	
Reviewed by	Moderator: Reviewers:		
Approved by	AMB & PMB		

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	30-03-2012	First draft	T. Ferrari/EGI.eu G. Borges/LIP P. Solagna/EGI.eu
2			
3			

IV. APPLICATION AREA

This document is a formal deliverable for the European Commission, applicable to all members of the EGI-InSPIRE project, beneficiaries and Joint Research Unit members, as well as its collaborating projects.

V. DOCUMENT AMENDMENT PROCEDURE

Amendments, comments and suggestions should be sent to the authors. The procedures documented in the EGI-InSPIRE “Document Management Procedure” will be followed:

<https://wiki.egi.eu/wiki/Procedures>

VI. TERMINOLOGY



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

VII. INDEX OF FIGURES

Figure 1. Number of certified production RCs from 01/09/2009 to 01/03/2012 (data source: GOCDB).	14
Figure 2. EGI countries hosting certified production Resource Centres from 01/09/2009 to 01/03/2012 (data source: GOCDB).	16
Figure 3. Total number of CPU cores provided by EGI RPs from PQ1 (June 2010) to PQ7 (January 2012) by RPs who are EGI-InSPIRE partners or EGI Council members (blue bar), and the amount of resources contributed together with integrated infrastructures (red bar).	20
Figure 4. Distribution of logical cores (blue bar) and HEP SPEC 06 installed capacity (red bar) at the end of PQ7 across EGI Resource infrastructure Providers, including EGI-InSPIRE partners, council members and integrated infrastructures.	20
Figure 5. Installed disk capacity in PB across the EGI RPs at the end of PQ7 – red bar – compared to the installed capacity in PQ3 – blue bar (source: Metrics Portal and Gstat).	21
Figure 6. Distribution of number Vos per discipline (March 2012, source: Operations Portal).	24
Figure 7. Comparison of the VO distribution at the end of March 2011 (blue bars) and at the end of March 2012 (red bars).	24
Figure 8. User distribution per discipline (March 2012, source: Operations Portal)	25
Figure 9. Comparison of the number of users per discipline in April 2011 (blue bar) and April 2012 (red bar).	25
Figure 10. Usage of EGI resources (HEP-SPEC 06 CPU wall clock hours) from the beginning of the project to date - a – and during PY2 - b (source: accounting portal).	27
Figure 11. Used normalized CPU wall clock time (a) and number of jobs done (b) across disciplines during PY2.	29
Figure 12. Distribution across EGI Operations Centres of usage of CPU wall clock time (HEP-SPEC 06 hours). HEP usage is displayed in blue and the aggregated usage of non-HEP disciplines is in red from May 2011 to April 2012 (source: accounting portal).	31
Figure 13. Distribution across EGI Operations Centres of aggregated usage of non-HEP disciplines (CPU wall clock time in HEP-SPEC 06 hours) from May 2011 to April 2012 (source: accounting portal).	32
Figure 14. Distribution of resource usage (%) across HEP and non-HEP disciplines. Usage is estimated as normalized CPU wall clock time (HEP-SPEC 06 hours) from May 2011 to April 2012 (source: accounting portal).	32
Figure 15. Classification of VO resource utilization per activity level (Low – red, Medium – Green, High – blue) from May 2008 (beginning of EGEE-III project to March 2012).	33
Figure 16. Quarterly availability and reliability of resource centres averaged across EGI from May 2010 to end of PQ7.	36
Figure 17. Classification of NGIs by the monthly availability provided by their top-BDII services. Here NGIs are classified according to the average monthly availability calculated during the reference period September 2011 – January 2012 (end of PQ7).	37



Figure 18. ROD average monthly workload (from PQ5 to PQ7)..... 39

Figure 19. ROD quality metric (monthly average from PQ5 to PQ7)..... 39

Figure 20, deployment of the four reference grid middleware stacks across the EGI-InSPIRE operations centres, March 2012 (source GOCDB)..... 44

Figure 21, Number of instances of the different implementations of the compute capability, across the EGI_InSPIRE partners and the integrated Resource infrastructure Providers, March 2012 (source: GOCDB)..... 45

Figure 22. Deployment of the different (gLite and EMI) releases for the services: WN, CREAM, WMS, LFC and DPM, March 2012 (source: information discovery system). 47

Figure 23, Number of core services instances deployed within the EGI-InSPIRE integrated infrastructure. Where not specified, source: information system (March 2012)..... 50

Figure 24. Number of Staged Rollout tests performed from PQ1 to PQ8 by NGIs/EIROs. 52



VIII. 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.



IX. EXECUTIVE SUMMARY

The total number of Resource Centres (RCs) in March 2011 amounts to 352 instances, of which: 326 are contributed by European NGIs/EIROs that are EGI-InSPIRE partners or Council members and 26 by integrated Resource infrastructure Provides (RPs) namely: Canada, Latin America – Brazil, Chile, Colombia and Mexico, and Iniciativa de Grid de America Latina - Caribe – Argentina, Brazil, Chile, Venezuela). Of these 326 RCs, 27 are contributed by Asia Pacific NGIs. The number of Resource Centres has been increasing at a lower speed since PQ2: while PY1 was characterized by the expansion of the infrastructure in the Baltic region, Latin America and Asia Pacific, PY2 was stable, and in PY3 further expansions will be mainly driven by the integration of new providers or by the integration of new software platforms.

In PY2 two new RPs got engaged with EGI by signing a Resource infrastructure Provider MoU: the South African Grid Initiative and the Ukrainian National Grid. In March 2012 EGI comprises resources provided across 56 countries and one European Intergovernmental Research Institute - CERN. 12 countries contribute resources through Resource Infrastructure Providers that are non-EGI-InSPIRE partners but are fully integrated with the EGI Services Infrastructure.

EGI currently comprehends 27 national operations centres and 9 federated operations centres encompassing multiple NGIs. The existing federated centres in Europe (IberGrid, NGI_NL and NGI_IT) comprehend two countries and are the result of a collaboration agreement that is expected to continue in PY3. On the other hand, integrated federated centres in Asia Pacific and Latin America encompass a large number of countries. The creation of new national grid initiatives in those regions will depend on their trends of expansion and on national policies. A new operations centre was created in Finland (NGI_FI) in June 2011. Part of the Finnish resources is operated by NGI_NDGF and another set of sites by NGI_FI. Two new additional NGIs were rolled to production: NGI_IE in June 2011 and NGI_UK in November 2011.

At the end of PQ7 the total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members, amounted to 270,800 and 2.96 Million HEP-SPEC 06. The contributed compute capacity increase significantly: the number of CPU cores increased by 30.7% since March 2011, while the installed capacity in HEP-SPEC 06 increased by 49.5% thus largely exceeding the expected targets for PY2. Storage capacity also increased considerably: the total amount of reported installed disk capacity amounts to 139 PB (+31.4% yearly increase), while the tape capacity amounts to 134.3 PB (+50% yearly increase).

HPC integration activities focused on the objective of supporting loosely and tightly coupled multi-scale simulations across EGI and PRACE resources. Pilot activities were carried out in collaboration with the MAPPER project. A joint MAPPER-EGI-PRACE task force helped with the operational integration of the MAPPER software platform.

The overall number of international and national VOs registered in the Operations Portal at the end of March 2012 amounts to 226 (3.20% from March 2011), including 20883 registered users (14.30% increase from March 2011). The disciplines which recorded the larger number of VOs during PY2 are in decreasing order: High-Energy Physics (+23.0% , with +39.25% of new users registered), Others (+17.7, with +19.41% of new users), Multidisciplinary VOs (+17.% VOs and +15.26% users), Infrastructure (+12.8 VOs and +2.95% users), Astronomy Astrophysics and Astro-particle Physics (+11.9% VOs and +12.51% users), Life Sciences (+8% VOs and + 5.61% users), Earth Sciences (+5.3% VOs and +1.55% users), Computer Science and Mathematics (+6% VOs and +0.20% users), Computational Chemistry (+2.2 VOs and + 2.91% users) and Fusion (0.9% VOs and +0.46% users).



The PY2 overall quantity of computing resources used amounts to 10.5 Billion HEP-SPEC06 Hours (the corresponding amount of consumed resources consumed during PY1 is 6.8 Billion HEP-SPEC06 Hours). The PY2 workload was generated by 492.5 Million jobs (1.35 Million Job/day on average).

Overall resource utilization has been satisfactorily progressing confirming the trends of PY1. The yearly increase of the total number of jobs executed in the infrastructure in the period May 2011-April 2012 amounts to +46.42% of the yearly job workload done from May 2010 to April 2011. The increase is even higher if looking to the CPU time consumed: the normalized CPU wallclock time (HEP-SPEC06 hours) from May 2011 to April 2012 amounts to +52.91% of the same amount during May 2010 – April 2011.

The High-Energy Physics discipline (contributing 39.25% of the user community) is still expanding in resource utilization, and the used normalized CPU wall time increased from 91.13% (April 2011) to 93.60% (April 2012) of the overall amount of EGI resources used. In particular, the used normalized CPU wall time yearly increased by +48.82%, while the job rate yearly increased by +57.06%. Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.25% of the overall EGI used CPU wall clock time, showing a +117.79% yearly increase from April 2011. Life Sciences are the third community for usage (1.30% of the overall EGI used normalized CPU time).

Usage distribution naturally reflects availability of installed capacity; however the level of multidisciplinary support varies considerably across the infrastructures. NGI_IT is the infrastructure with the largest absolute amount of resources used by non-HEP communities (more than 128 Million hours used). HEP usage is dominant in large resource infrastructures, while support of other disciplines dominates in Latin America and several countries in Eastern Europe, where in some cases it equals to 100% of the entire amount of used resources.

Despite of the large number of registered VOs, the number of “active” ones is a significantly smaller fraction. At the end of PQ7, 56 active VOs were reported by the Accounting Portal, of which 8 featuring low activity, 23 medium and 25 high. A short-term increase in the number of active VOs was observed starting in October 2009. This increase, terminated in April 2010, was followed by a progressive decrease that affected the Low Activity category, which presumably includes new user communities in their initial pilot activities. In order to address this problem, the long-term EGI strategy aims at extending the current user base by facilitating the deployment of heterogeneous application platforms and by embracing cloud provisioning models.

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 – March 2012 (second year of EGI-InSPIRE). Several NGIs already integrated in PY1, were affected by periodic performance instability because of lack of expertise in technical services deployed. A support action was kicked off in PQ6 in collaboration with the Greek JRU to technically support those infrastructures.

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 now reported monthly. Purpose of this reporting is to check the availability and reliability of core services operated by NGIs and EIROs, which is typically highly critical as these services provide access to RC services.

The EGI VO Services – defined at the beginning of EGI-InSPIRE as part of NA3 and migrated to SA1 during PY2 – aim at supporting VOs in the whole process of start-up, management and operation, pointing out to tools, services, documentation and guidelines to maximize the usage of the resources,



easing service deployment, and bridging the VO community with the infrastructure need. VO services are mature enough to be supported by NGI operational teams and the expertise on operating those services is also widely available in the operations community.

The number of operated software platforms was increased from three at the end of PY1 to five in PY2. Currently the grid middleware stacks fully deployed in the infrastructure are: ARC, gLite, GLOBUS, and UNICORE. gLite is the mostly deployed middleware, but the number of operations centres supporting non-gLite stacks slightly increased during PY2. A gLite 3.1 end of life plan also including a few other products, was defined in February 2012. The support calendar of gLite 3.2 – which reached end of life in April 2012 – will be discussed in May 2012. The update of the EMI 1 release depends on the service and the new features released with the products.

The EGI integrated infrastructure comprehends 359 core services: 64 VOMS instances, 150 WMS, 41 LFC and 104 Top-BDII.

The Staged Rollout activities performed by SA1 consolidated in PY2 and are now clearly perceived as EGI added value to the software provided by mainstream technology providers. The largest number of products was tested in PQ5 in preparation to the release of UMD 1.0 (81 tests in total resulting in 2 products rejected). This number was gradually reduced in the following quarters, as subsequent UMD updates only included a subset of products being updated. The number of participating Early Adopters has been progressively increasing to test a growing set of products from EMI, IGE and EGI-InSPIRE JRA1 (operational tools). Staged Rollout is a community effort currently contributed by 56 distributed teams.



TABLE OF CONTENTS

1	INTRODUCTION	11
2	RESOURCE INFRASTRUCTURE	13
2.1	Resource Centres	13
2.2	Resource infrastructure Providers	14
2.3	Operations Centres	18
3	INSTALLED CAPACITY	19
3.1	Compute Resources	19
3.2	Storage Resources	21
3.3	Compute Resources for Parallel Jobs	21
4	DISCIPLINES, VIRTUAL ORGANIZATIONS AND USERS	23
4.1	VO Distribution across scientific fields	23
4.2	User Distribution across scientific fields	24
4.3	Resource Utilization per Discipline	26
5	RESOURCE UTILIZATION	30
5.1	Active Vos	33
6	SERVICE LEVELS	34
6.1	Service Level Targets and Reporting	34
6.2	RC Performance	34
6.2.1	Availability and Reliability	34
6.3	RP Performance	36
6.3.1	Availability and Reliability	37
6.3.2	ROD Performance Index	38
6.4	NGI Support Actions	40
7	VO SERVICES	41
7.1	Status	41
7.2	Utilization	42
8	GRID SERVICES	43
8.1	Integrated Software Platforms	43
8.2	Deployment of gLite releases	45
8.3	Core Middleware Services	47
9	STAGED ROLLOUT INFRASTRUCTURE	51
10	CONCLUSIONS AND FUTURE WORK	53
11	REFERENCES	55
12	ANNEX I. VO DISTRIBUTION PER DISCIPLINE	57
12.1	Astronomy Astrophysics and Astro-particle Physics	57
12.2	Computer Science and Mathematics	57
12.3	Earth Sciences	58



12.4	Fusion.....	58
12.5	High Energy Physics	59
12.6	Infrastructure	59
12.7	Life Sciences.....	60
12.8	Multidisciplinary VOs.....	60
12.9	Other Disciplines.....	61

1 INTRODUCTION

This document illustrates the status and progress of the EGI production infrastructure and of the Staged Rollout infrastructure at the end of PY2. The current status of Resource Centres (RCs), Resource infrastructure Providers (RPs) and of the Operations Centres is shown in 2. The amount of installed capacity being operated, the status and trends of the existing EGI user community, 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 and the RP level, and analyses the trends in performance delivered and the actions being implemented to support a number of infrastructures requiring consolidation. Sections 7 and 8 describe the status of VO-specific services and of general core infrastructure services. The Staged Rollout infrastructure is described in Section 9. Section 10 summarises the progress made in operations during PY2 to meet the project requirements and concludes the document.

This document is complemented by various milestones and deliverables, which provide additional information about SA1 services and activities:

- The EGI operations service catalogue and the set of operational procedures and activities are described in the EGI Operations Architecture: Grid Service Management Best Practices [D4.3].
- The central operational tools provided by EGI.eu and the status assessment of EGI operational Global Tasks, are not part of this deliverable. Information is provided instead by the EGI Global Task Review [MS115], while information about the related development activities is available in the Annual Report on Operational Tool Maintenance [D7.2]. Feedback about the operations global services provided by EGI, was gathered from NGIs in January and results are available on-line¹.
- Progress of NCI operational activities is documented in the NCI International Task Review [MS116].
- Security in the EGI infrastructure, its scope and aims, and the assets that EGI security seeks to protect are described in the Security Risk Assessment of the EGI Infrastructure [D4.4]. The work of the various security groups in or associated with EGI is briefly described there, together with practices and standards for IT security and their usage and possible future. Some security incidents have occurred over the last year, these are also briefly documented including how they were handled. Previous Overall Security Risk assessments are then summarized, and plans for a security risk assessment, which is taking place during PQ8 and PQ9, are described.
- The evolution of the EGI Operational Service Agreement framework including the reporting system needed to produce monthly reports is detailed in the Operational Level Agreements (OLAs) within the EGI production infrastructure [MS418].

¹ <https://www.egi.eu/indico/getFile.py/access?contribId=3&resId=0&materialId=1&confId=618>



- Technical information about progress of the operational integration of new software stacks is provided in milestone Integrating Resources into the EGI Production Infrastructure [MS414] and on wiki².

² <https://wiki.egi.eu/wiki/Interoperations>

2 RESOURCE INFRASTRUCTURE

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 are described in this section.

2.1 Resource Centres

The Resource Centre is defined to be 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 2012).

Resource Centres	Number of RCs
EGI-InSPIRE Partners and NGI Council Members/PY2 Target	326/330
From non-European EGI-InSPIRE Partners	27
From integrated Infrastructures (Canada, Latin America, IGALC)	26
Total	352

As shown in the table above, the total number of RCs in March 2011 amounts to 352 instances, of which: 326 are contributed by European NGIs/EIROs that are EGI-InSPIRE partners or Council members and 26 by integrated RPs namely: Canada, Latin America – Brazil, Chile, Colombia and Mexico, and *Iniciativa de Grid de America Latina – Caribe* – Argentina, Brazil, Chile, Venezuela).

Of the 326 RCs mentioned above, 27 are contributed by Asia Pacific NGIs.

The target for PY2 of 330 RCs was approximately met (326 RCs certified centres against 330 expected), scoring a -1.21% deviation. Small fluctuations are normal, as centres can be temporarily suspended for various different reasons (for example in case of low performance, security problems or infrastructure upgrades)³. Every month an average number of 3-5 centres are eligible to suspension because of insufficient performance.

The total amount increased from 346 (PQ5) to 348 (PQ6), and 347 in PQ7. Figure 1 plots the number of RCs from January 2009, and it shows that the number has been increasing at a lower speed since PQ2. In fact, PY1 was characterized by the expansion of the infrastructure in the Baltic region (Bosnia and Herzegovina, South East Europe region (Armenia, FYR of Montenegro and Georgia), Latin America (Argentina, Venezuela) and Asia Pacific (New Zealand).

From PY3 further expansions will be driven mainly driven by the integration of new providers and the capability of integrating new software platforms. In PY2 two new RPs got engaged with EGI by signing a Resource infrastructure Provider MoU: the South African Grid Initiative [SAG] and the Ukrainian National Grid [UNG]. An expansion of the number of RCs is expected in PY3 in those two areas.

GOCDB was used to extract information about the number of certified production RCs.

³ https://wiki.egi.eu/wiki/Underperforming_sites_and_suspensions




Figure 1. Number of certified production RCs from 01/09/2009 to 01/03/2012 (data source: GOCDB).

2.2 Resource infrastructure Providers

The Resource infrastructure Provider is the legal organisation responsible for any matter that concerns the 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 at an international level. Examples of a Resource infrastructure Provider are the European Intergovernmental Research Organisations (EIRO) and the National Grid Initiatives (NGIs) [GLO].

In March 2012 EGI comprises resources provided across 56 countries and one European Intergovernmental Research Institute – CERN.

In 42 countries resources are directly contributed by EGI-InSPIRE partners or RPs that are Council members. On the other hand, 12 countries contribute resources through Resource Infrastructure Providers that are non-EGI-InSPIRE partners but are fully integrated with the EGI Services Infrastructure.

- India, Pakistan and New Zealand (Asia Pacific Federation);
- Austria (Italian Federation): operations support of Austrian RCs moved from NGI_NDGF to NGI_IT during PY2;

- Argentina, Brazil, Chile, and Venezuela (IGALC);
- Canada (Canada Federation);
- Brazil, Chile, Colombia and Mexico (Latin America Federation);
- Ukraine (Russian Federation).

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

The number of countries contributing resources is approximately constant. Since PQ5 the production infrastructure in the Philippines has been instable, and RCs have been periodically removed from the infrastructure. The growth trend since 01 September 2009 is illustrated in **Figure 2**.

In March 2012, EGI-InSPIRE partners in four countries are not contributing resources: Albania, Indonesia, Philippines and Singapore. The integration of Moldova was completed in March 2012.




Figure 2. EGI countries hosting certified production Resource Centres from 01/09/2009 to 01/03/2012 (data source: GOCDB).

Table 2.

The number of countries contributing resources is approximately constant. Since PQ5 the production infrastructure in the Philippines has been unstable, and RCs have been periodically removed from the infrastructure. The growth trend since 01 September 2009 is illustrated in **Figure 2**.

In March 2012, EGI-InSPIRE partners in four countries are not contributing resources: Albania, Indonesia, Philippines and Singapore. The integration of Moldova was completed in March 2012.




Figure 2. EGI countries hosting certified production Resource Centres from 01/09/2009 to 01/03/2012 (data source: GOCDB).

Table 2. Distribution of production certified RCs across countries and Operations Centres in PQ5, PQ6, PQ7 and March 2012. Philippines currently do not contribute resources. Data source: GOCDB.

I = Integrated, EP = European Partner, P = non-European partner, C = only Council member

Operations Centre	Country		PQ4	PQ5	PQ6	PQ7	01/03/2012
IGALC	Argentina	I	1	1	1	1	1
NGI_ARMGRID	Armenia	EP	4	1	1	1	3
Asia Pacific	Australia	P	2	2	1	1	1
NGI_IT	Austria	I	2	2	2	2	2
NGI_BY	Belarus	EP	5	5	4	4	4
NIG_NL	Belgium	C	4	4	3	3	3
NGI_BA	Bosnia and H.	EP	1	1	1	1	1
IGALC/Latin America	Brazil	I	4	4	4	5	5
NGI_BG	Bulgaria	EP	9	9	9	9	9
Canada	Canada	I	7	7	7	7	7
IGALC/Latin America	Chile	I	1	1	1	2	2
Asia Pacific/Canada	China	P	2	2	3	3	3
Latin America	Colombia	I	1	1	1	1	1
NIG_HR	Croatia	EP	3	3	3	3	3
NGI_CY	Cyprus	EP	2	2	2	2	2
NGI_CZ	Czech Republic	EP	3	3	2	2	2
NGI_NDGF	Denmark	EP	2	2	2	2	2
NGI_NDGF	Estonia	C	2	2	2	2	2
NGI_FI/NGI_NDGF	Finland	EP	1	2	2	2	2
NGI_FRANCE	France	EP	20	20	19	19	19
NGI_MARGI	FYROM	EP	2	2	2	2	2
NGI_GE	Georgia	EP	1	1	1	1	1
NGI_DE	Germany	EP	19	21	21	21	21
NGI_GRNET	Greece	EP	14	14	14	15	15
NGI_HU	Hungary	EP	5	5	4	4	4
Asia Pacific	India	I	2	2	2	2	2
NGI_IE	Ireland	EP	6	6	6	6	6
NGI_IL	Israel	EP	3	3	3	4	6
NGI_IT	Italy	EP	56	56	58	61	61
Asia Pacific	Japan	P	4	3	3	3	3
NGI_NDGF	Latvia	EP	3	2	2	2	2
NGI_NDGF	Lithuania	EP	3	4	2	2	2
Asia Pacific	Malaysia	P	4	4	4	4	5
Latin America	Mexico	I	1	1	2	2	2
NGI_ME	Montenegro	EP	1	1	1	1	1
NGI_NL	Netherlands	EP	15	15	16	16	16
Asia Pacific	New Zealand	I	1	1	1	1	1
NGI_NDGF	Norway	EP	1	1	1	1	1
Asia Pacific	Pakistan	I	1	1	1	1	2
Asia Pacific	Philippines	P	1	1	0	0	0
NGI_PL	Poland	EP	8	8	8	9	10
NGI_IberGrid	Portugal	EP	8	8	7	7	7
NGI_RO	Romania	EP	9	11	11	11	12
Russia	Russia	EP	12	12	11	10	10
NGI_AEGIS	Serbia	EP	6	6	6	6	6
NGI_SK	Slovakia	EP	4	4	4	4	4
NGI_SI	Slovenia	EP	2	2	2	2	2
Asia Pacific	South Korea	P	2	3	4	4	4
NGI_IberGrid	Spain	EP	21	22	22	24	25
NGI_NDGF	Sweden	EP	1	1	1	1	1
NGI_CH/CERN	Switzerland	EP	8	8	8	7	6
Asia Pacific	Taiwan	P	6	6	5	5	6
Asia Pacific	Thailand	P	2	2	1	2	2
NGI_TR	Turkey	EP	6	6	6	5	4
Russia	Ukraine	P	1	2	3	3	3
NGI_UK	United Kingdom	EP	21	21	21	22	22
IGALC	Venezuela	I	1	1	1	1	1
TOT Resource Centres			338	342	335	344	352
TOT Countries			57	57	56	56	56

2.3 Operations Centres

The Operations Centre is defined to be a centre offering operations services on behalf of the Resource Infrastructure Provider [GLO].

EGI currently comprehends 27 national operations centres and 9 federated operations centres encompassing multiple NGIs (Table 3). The existing federated centres in Europe (IberGrid, NGI_NL and NGI_IT) comprehend two countries and are the result of a collaboration agreement that is expected to continue in PY3. On the other hand, integrated federated centres in Asia Pacific and Latin America encompass a large number of countries. The creation of new national grid initiatives in those regions will depend on their trends of expansion and on national policies.

A new operations centre was created in Finland (NGI_FI) in June 2011. Part of the Finnish resources is operated by NGI_NDGF and another set of sites by NGI_FI.

Two new additional NGIs were rolled to production: NGI_IE in June 2011 and NGI_UK in November 2011. Both operations centres were created in order to proceed with the decommissioning of an EGEE legacy operations centre providing operational services in both countries (UKI). In November 2011 and additional EGEE legacy operations centre (DECH) was decommissioned.

Finally, in November 2011 the operations of Austrian sites were handed off by NGI_NDGF to NGI_IT. Two new Resource infrastructure Provider MoUs were agreed in PY2, and these will result in the certification of two new operations centres (NGI_UA and NGI_ZA) which are expected to be certified in PY3. The transition of EGEE legacy ROCs to NGIs was completed in PY2.

Table 3. List of EGI federated Operations Centres

Federated operations centres	Member countries	Comments
Asia Pacific	Australia, China, India, Japan, Malaysia, New Zealand, Pakistan, South Korea, Taiwan, Thailand	Philippines is affiliated to the Asia Pacific operations centre, the status of site PH-ASTI-LIKNAYAN has been periodically transitioning from uncertified to certified status. Viet namese sites were decommissioned during PY1 (August 2010) and never taken back to production afterwards.
Canada	Canada, China	
IberGrid	Portugal, Spain	
Iniciativa de Grid de America Latina – Caribe (ROC_IGALC)	Argentina, Brazil, Chile, Venezuela	Argentina whose site was certified in PQ2, has been stably running in production since
Latin America (ROC_LA)	Brazil, Chile, Colombia, Mexico	
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, Finland, Latvia, Lithuania, Norway, Sweden	
Netherlands (NGI_NL)	Belgium, Netherlands	
Russia (Russia)	Russia, Ukraine	Migration of Ukrainian sites expected in PY3 after the consolidation of a new national operations centre in Ukraine



3 INSTALLED CAPACITY

Installed capacity is monitored at the end of each project quarter. 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 published.

ARC, GLOBUS and UNICORE are currently not published in the Information Discovery System with the exception of ARC resources operated by NGI_NDGF. EGI is collaborating with both the EMI and IGE project for the development of a unified Information Discovery System.

3.1 Compute Resources

The total amount of CPU cores contributed by EGI-InSPIRE partners and RPs council members amounts to 270,800 and 2.96 Million HEP-SPEC 06 at the end of PQ7, while the total number including compute resources contributed by integrated and peer infrastructures amounts to 399,300 unit. This value abundantly exceeds the PY2 target of 350,000 total cores.

Looking at the compute resources provided by EGI partners (EGI-InSPIRE partners and council members) – Table 5, the number of CPU cores increased by 30.7% since March 2011, while the installed capacity in HEP-SPEC 06 increased by 49.5%. This significant increase reflects the advancements in CPU technology and core density per CPU and shows a trend in building up capacity in the existing RCs.

Table 4. EGI-InSPIRE logical CPUs

Logical CPUs	PQ7/PY2 Target
EGI partners	270,800
EGI-InSPIRE partners plus integrated and peer infrastructures	399,300/350,000

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

	April 2010 (EGEE-III Infrastructure)	March 2011 (EGI-InSPIRE Infrastructure)	March 2012 (EGI-InSPIRE Infrastructure)
Logical CPUs/yearly increase	192,000	207,203/+7.9%	270,800/+30.7%
Million SI00	335 (1.34 Million HEP-SPEC 06)	495/+47.7% (1.98 Million HEP-SPEC 06)	740/+49.5% (2.96 Million HEP-SPEC 06)




Figure 3. Total number of CPU cores provided by EGI RPs from PQ1 (June 2010) to PQ7 (January 2012) by RPs who are EGI-InSPIRE partners or EGI Council members (blue bar), and the amount of resources contributed together with integrated infrastructures (red bar).




Figure 4. Distribution of logical cores (blue bar) and HEP SPEC 06 installed capacity (red bar) at the end of PQ7 across EGI Resource infrastructure Providers, including EGI-InSPIRE partners, council members and integrated infrastructures⁴.

⁴ Denmark and ROC Latin America are not included in the diagram because of errors affecting the installed capacity values published in the information discovery system.

3.2 Storage Resources

Information about storage capacity provided by the EGI-InSPIRE partners is periodically collected by the Metrics Portal from the Information Discovery System (Gstat) and validated by EGI-InSPIRE partners. As already mentioned for the compute capacity, accuracy of information 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 sites.

At the end of PQ7, the total amount of reported installed disk capacity amounts to 139 PB (+31.4% 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 seven NGIs/EIROs, which are in descending order: Germany, CERN, United Kingdom, Italy, France, Spain and The Netherlands.




Figure 5. Installed disk capacity in PB across the EGI RPs at the end of PQ7 – red bar – compared to the installed capacity in PQ3 – blue bar (source: Metrics Portal and Gstat).

Tape capacity is mainly provided by CERN and WLCG Tier-1 RCs. At the end of PQ7 the total installed tape (also known as *nearline*) capacity reported in Gstat amounts to 134.3 PB (+50% yearly increase).

3.3 Compute Resources for Parallel Jobs

Information is gathered periodically in the project quarterly reports about the number of *high-performance* clusters operated. With high-performance we refer to clusters that feature a local high-

speed low-latency interconnect (e.g. Myrinet⁵, InfiniBand⁶). The clusters that qualify as high-performance, have been reported by the Resource Infrastructure Providers to be slightly less than at the end of PQ3. They amount in total to 40 units (-25.9% yearly decrease) at the end of PQ7.

Table 6. Integration metrics (HPC and MPI)

Metric	PQ3	PQ7/PY2 Target
Number of HPC clusters (M.SA1.Integration.1)	54	40/3
Number of sites with MPI (M.SA1.Integration.2)	90	108/100

At the end of PQ3 Message Passing Interface [MPI] jobs were supported by 108 sites (+20% yearly increase) as shown in Table 6.

During PY2, HPC integration activities focused on the objective of supporting loosely and tightly coupled multi-scale simulations across EGI and PRACE resources. Pilot activities were carried out in collaboration with the MAPPER project⁷. A joint MAPPER-EGI-PRACE task force⁸ helped with the operational integration of the MAPPER software platform. The objectives of the task force are;

- to integrate QCG Nagios probes into the SAM release and to assess which of these probes will affect site Availability/Reliability in case of failure (affecting
- to assess the status of accounting, and the developments needed to integrate what exists with the EGI accounting infrastructure
- to integrate QCG support activities into the EGI mainstream support tasks (1st level, 2nd level and 3rd level).

Integration with GOCDB and SAM are now almost completely accomplished. Accounting integration is being developed so that usage records for EGI resources can be accounted centrally. MAPPER integration requirements will be evaluated to explore the integration across a selected set of EGI and PRACE operational services.

⁵ <http://www.myri.com/myrinet/overview/>

⁶ <http://www.infinibandta.org/>

⁷ www.mapper-project.eu/

⁸ https://wiki.egi.eu/wiki/MAPPER-PRACE-EGI_Task_Force_%28MTF%29-II



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.

The overall number of international and national Vos registered in the Operations Portal⁹ at the end of March 2012 amounts to 226 (3.20% from March 2011), including 20883 registered users (14.30% increase from March 2011).

Currently user statistics extracted from VO Membership Services do not provide information about users suspended because of expired membership, and about the number of active users being identified by robot certificates. These deficiencies affect the accuracy of statistics herein reported. VOMS extensions have been requested in PY2, and will be released starting with the EMI 2.0 VOMS version to be released in May 2012¹⁰.

4.1 VO Distribution across scientific fields

The disciplines which recorded the larger number of Vos during PY2 are in decreasing order: High-Energy Physics (+23.0% , with +39.25% of new users registered), Others (+17.7, with +19.41% of new users), Multidisciplinary Vos (+17.7% Vos and +15.26% users), Infrastructure (+12.8 Vos and +2.95% users), Astronomy Astrophysics and Astro-particle Physics (+11.9% Vos and +12.51% users), Life Sciences (+8% Vos and + 5.61% users), Earth Sciences (+5.3% Vos and +1.55% users), Computer Science and Mathematics (+6% Vos and +0.20% users), Computational Chemistry (+2.2 Vos and + 2.91% users) and Fusion (0.9% Vos and +0.46% users).

The distribution of Vos per discipline is illustrated in Figure 7. Disciplines that have increased their relative share are High Energy Physics (23.0%), Multidisciplinary Vos (16.8%) and Astronomy Astrophysics and Astro-particles Physics (11.9%).

⁹ <http://operations-portal.egi.eu/vos>

¹⁰ <http://www.eu-emi.eu/emi-2-matterhorn>

VO distribution (%) per discipline




Figure 6. Distribution of number VOs per discipline (March 2012, source: Operations Portal).

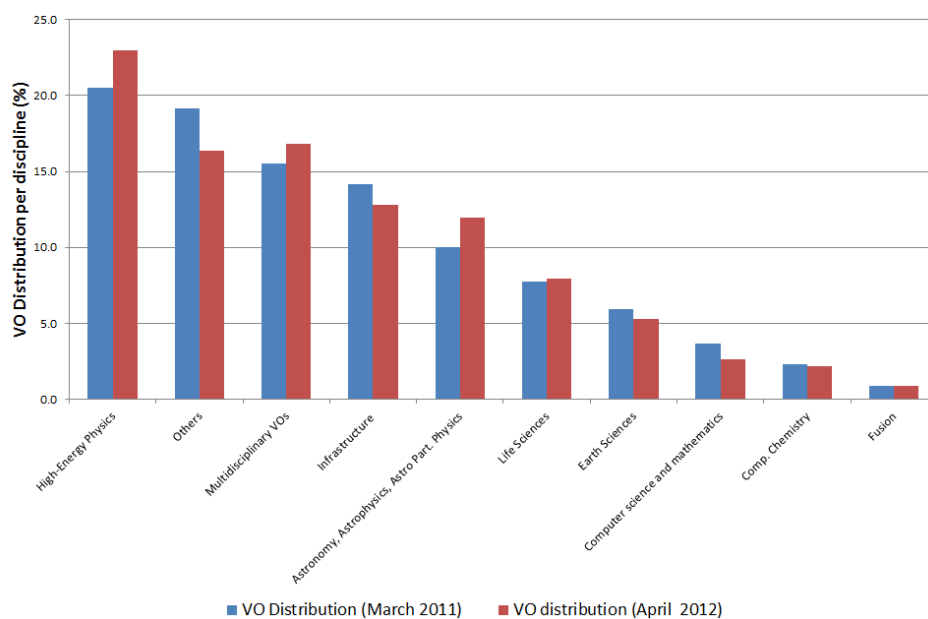


Figure 7. Comparison of the VO distribution at the end of March 2011 (blue bars) and at the end of March 2012 (red bars).

4.2 User Distribution across scientific fields

As illustrated in the largest disciplines in terms of number of registered users are: High Energy Physics (39.25%), Others (19.41%) and Multidisciplinary VOs (15.26%).

During PY2 the number of registered users for several disciplines has increased considerably, most notably Astronomy Astrophysics and Astro-particle Physics (+420.28%) with the most active VO being AUGER (66.14% of the total amount of normalized CPU time used in the discipline) and

Computational Chemistry (+216.15) with COMPCHEM being the most active VO (89.36% of normalized CPU time used in the discipline). The detailed user distribution per discipline is presented in Appendix.




Figure 8. User distribution per discipline (March 2012, source: Operations Portal)




Figure 9. Comparison of the number of users per discipline in April 2011 (blue bar) and April 2012 (red bar).

4.3 Resource Utilization per Discipline

Overall resource utilization during PY2 has been progressively increasing confirming the trend of PY1. The yearly increase of the total number of jobs executed in the infrastructure in the period May 2011-April 2012 amounts to +46.42% of the yearly job workload done from May 2010 to April 2011.

The increase is even higher if looking to the CPU time consumed: the normalized CPU wallclock time (HEP-SPEC06 hours) from May 2011 to April 2012 amounts to +52.91% of the same amount computed from May 2010 to April 2011.

HEP-SPEC06 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-SPEC06 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 1K.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 interrogates the site 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 and during PY2 are both shown in Figure 10.

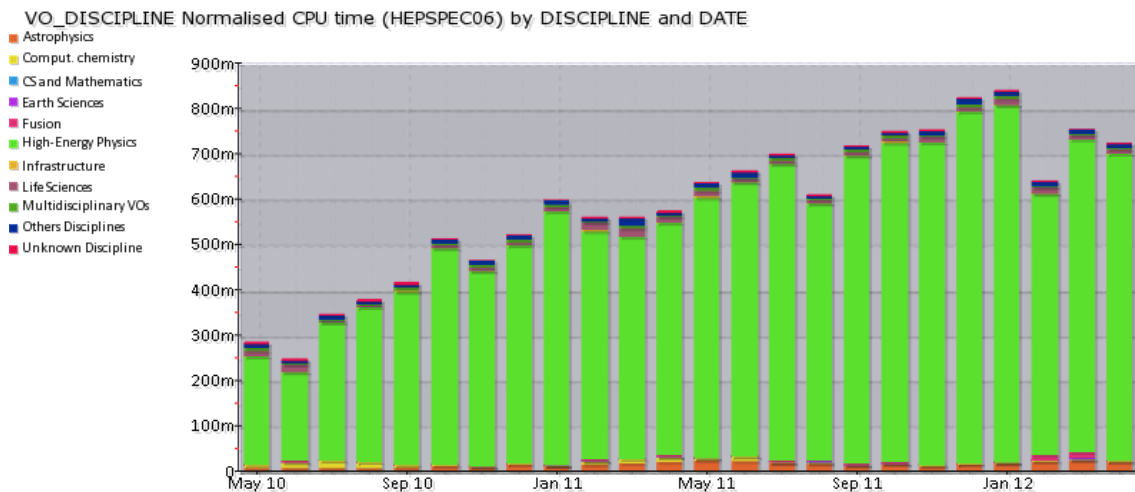
The High-Energy Physics discipline (contributing 39.25% of the user community) is still expanding in resource utilization, and its used normalized CPU wall time increased from 91.13% (April 2011) to 93.60% (April 2012) of the overall EGI usage amount – see Table 7. In particular, the used normalized CPU wall time yearly increase amounts to +48.82%, while the job rate yearly increase is +57.06%.

As we can see from Table 7, while the relative distribution of used CPU wall clock time across the disciplines at the end of PY2 (column A) is approximately equal to what recorded at the end of PY1 (column C), because of the major increase in the overall amount of installed capacity, several disciplines increased their usage of the infrastructure.

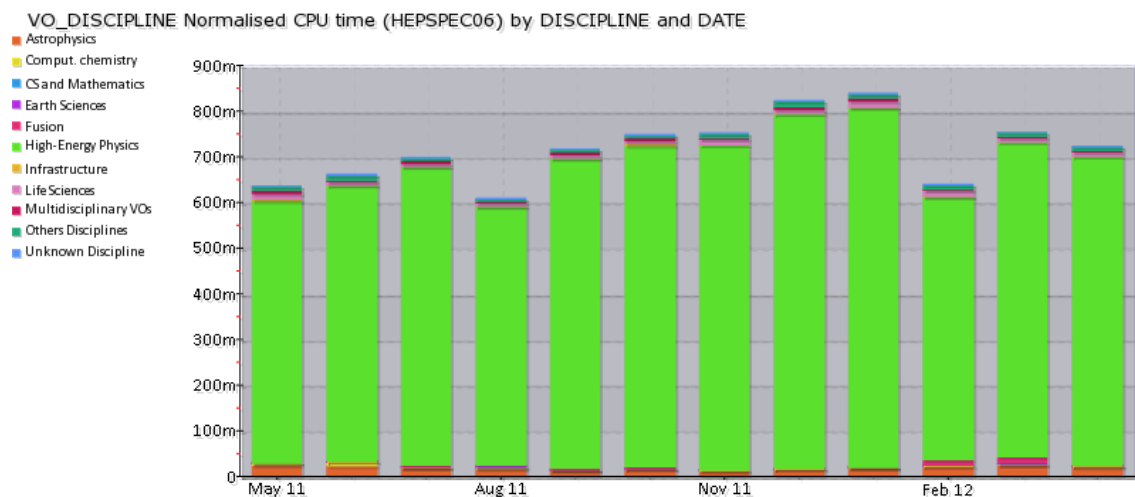
Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.25% of the overall EGI used CPU wall clock time, showing a +117.79% yearly increase from April 2011.

Life Sciences are the third community for usage (1.30% of the overall EGI used normalized CPU time). For this community the overall amount of used CPU wall clock time remained approximately the same (+1.97% yearly increase), while the job submission partner changed considerably as the rate of jobs done yearly increased by +42.54%.

“Other disciplines”, Fusion and Earth Sciences also increased both their used CPU wall clock time and the job rate. The overall trend of used normalized CPU wall clock time for non-HEP disciplines is plotted in Figure 11. 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 not proportional to the corresponding amount of CPU wall time consumed.



(a)



(b)


Figure 10. Usage of EGI resources (HEP-SPEC 06 CPU wall clock hours) from the beginning of the project to date - a – and during PY2 - b (source: accounting portal).

As shown in Figure 10, since the beginning of the project High-Energy Physics has greatly increased its workload in comparison with the one generated by the other disciplines.

Table 7. Used normalized CPU wall clock time and jobs done per discipline in PY1, PY2 and the respective early increase (source: accounting portal).

Discipline	May 11–April 12		May 10 – April 11		Jobs (yearly increase from May 10) €	CPU wall time (yearly increase from May 10) (F)
	% CPU n. wall time	% done Jobs	% CPU n. wall time	% of jobs done		
	(A)	(B)	(C)	(D)		
High-Energy Physics	93.60	91.58	91.13	90.10	+48.82%	+57.06%
Infrastructure	0.20	3.26	0.27	4.51	+5.91%	+13.44%
Life Sciences	1.30	1.75	1.94	1.80	+42.54%	+1.97%
Astrophysics	2.25	1.58	2.01	1.06	+117.79%	+71.13%
Multidisciplinary	0.39	0.48	0.83	0.86	-18.14%	-28.29%
Others Disciplines	1.23	0.72	1.55	0.70	+52.30	+20.86%
Unknown Discipline	0.20	0.29	0.86	0.63	-33.54	-63.89%
Comput. Chemistry	0.38	0.03	1.27	0.24	-23.51%	-54.45
Fusion	0.37	0.13	0.10	0.06	+238.31%	+454.15%
Earth Sciences	0.10	0.05	0.04	0.03	+118.31%	+260.29%
CS and Mathematics	0.00	0.03	0.00	0.00	+2303.19%	+28.99%

PRODUCTION Normalised Elapsed time (HEPSPEC06) by VO and DATE



© CESGA 'EGI View': PRODUCTION / normelap-HEPSPEC06 / 2011:5-2012:4 / VO-DATE / top10 (i) / ACCBAR-LIN / i 2012-04-30 16:12




Figure 11. Used normalized CPU wall clock time (a) and number of jobs done (b) across disciplines during PY2.

5 RESOURCE UTILIZATION

EGI accounting information is gathered and stored centrally for display through the accounting portal¹¹. Accounting information is aggregated by Operations Centre, whose list is obtained from GOCDDB.

Table 8. Annual compute resource usage (yearly figures)

	PY1	PY2/PY2 Target
Total normalized CPU wall clock time consumed (Billion HEP-SPEC 06 hours)	6.8	10.5
Jobs per year (Million)	334.8	492.5
Average number of Jobs per day (Million)	0.92	1.35/ 0.525

The overall quantity of computing resources used in PY2 amounts to 10.5 Billion HEP-SPEC06 Hours (the corresponding amount of consumed resources consumed during PY1 amounted to 6.8 Billion HEP-SPEC06 Hours) as shown in Table 8. The PY2 workload was generated by 492.5 Million jobs, which amounts to an average of 1.35 Million Job/day.

PY2 usage expressed in HEP-SPEC 06 Hours of CPU wall time across of the various resource infrastructures of EGI is plotted in Figure 12, 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 (red bars), the top infrastructures being (in decreasing order): NGI_UK, NGI_DE, NGI_France and NGI_IT. Usage distribution naturally reflects availability of installed capacity (Section 3), however the level of multidisciplinary support varies considerably across the infrastructures. Figure 13 plots the distribution of used HEP-SPEC 06 CPU wall clock hours of non-HEP user communities is distributed. NGI_IT is the infrastructure with the largest absolute amount of resources used by non-HEP communities with more than 128 Million hours.

The overall fraction of resources used by non-HEP communities in each infrastructure is plotted in Figure 14. The diagram shows how support of HEP is dominant in large resource infrastructures, while support of other disciplines dominates in Latin America and several countries in Eastern Europe where in some cases it equals 100% of the entire usage of resources. As outcome of the EGI-InSPIRE activities in outreach and technical support of new user communities introduced in PY2, it is expected that the fraction of usage non-HEP usage will increase in the coming project year.

¹¹ http://www4.egee.cesga.es/accounting/egee_view.php




Figure 12. Distribution across EGI Operations Centres of usage of CPU wall clock time (HEP-SPEC 06 hours). HEP usage is displayed in blue and the aggregated usage of non-HEP disciplines is in red from May 2011 to April 2012 (source: accounting portal).




Figure 13. Distribution across EGI Operations Centres of aggregated usage of non-HEP disciplines (CPU wall clock time in HEP-SPEC 06 hours) from May 2011 to April 2012 (source: accounting portal).




Figure 14. Distribution of resource usage (%) across HEP and non-HEP disciplines. Usage is estimated as normalized CPU wall clock time (HEP-SPEC 06 hours) from May 2011 to April 2012 (source: accounting portal).

5.1 Active Vos

Despite of the large number of registered Vos (226 according to the Operations Portal with a +3.20% of new Vos in PY2), the number of “active” Vos is a significantly smaller fraction. At the end of PQ7, 56 active Vos were reported by the Accounting Portal (metric M.NA3.8, [QR7]) of which 8 featuring low activity, 23 medium and 25 high. The Accounting Portal classifies Vos according to the following ranges:

- “Low Activity week”: CPU > 1 Day/Week, all the jobs run by the VO during the week consumed more than 24 hours (1 day) of CPU time, but less than 1 Month
- “Medium activity week”: CPU > 1 Month/Week, all the jobs run by the VO during the week consumed more than 1 month of CPU time and less than 1 year.
- “High activity week”: CPU > 1 Year/Week, all the jobs run by the VO during the week consumed more than 1 year of CPU time.




Figure 15. Classification of VO resource utilization per activity level (Low – red, Medium – Green, High – blue) from May 2008 (beginning of EGEE-III project to March 2012).

Figure 15 plots the long-term trend in the distribution between Low/Medium/High Activity Vos starting from the beginning of the EGEE-III project – May 2008 – to date. The diagram shows that the overall number of active Vos hasn’t significantly changed since 2008. A short-term increase trend started in October 2009 and finished in April 2010, was followed by progressive slow decrease that affected the Low Activity category (red), which presumably includes new user communities in their initial pilot activities.

The integration of heterogeneous grid middleware (ARC, dCache, gLite, GLOBUS and UNICORE) and of additional platforms, e.g. desktop grids and software platforms allowing the coupled usage of HTC and HPC resources across EGI and PRACE, is expected to lower the existing grid technical deployment barriers and to address a wider spectrum of user needs [INT] [MAP]. In addition, the long-term EGI strategy aims at extending the current user base by facilitating the deployment of heterogeneous application platforms by embracing cloud provisioning models [STR].

6 SERVICE LEVELS

Services are monitored at two different levels:

- Services provided at a RC level and
- Services provided at a RP level.

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 level and targets are formally defined in the RC Operational Level Agreement [RCO] and in the RP Operational Level Agreement [RPO].

6.1 Service Level Targets and Reporting

- Resource Centres¹²
 - Minimum Availability: 50%
 - Minimum Reliability: 50%
 - Reports:
https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics#Resource_Centres
- Resource infrastructure Providers¹³
 - 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#NGIs.2FEIROs

6.2 RC Performance

6.2.1 Availability and Reliability

Table 9. EGI-wide Availability and Reliability and the related project metric target.

EGI Average Monthly Reliability	May 2011-January 2012	Y2 Target
Reliability	95.42 %	91%
Availability	94.50 %	-

¹² https://wiki.egi.eu/wiki/SLM/RC_Service_Levels

¹³ https://wiki.egi.eu/wiki/SLM/RP_Service_Levels



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 ($[\text{uptime} / \text{total time}] * 100$), 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 ($[\text{uptime} / (\text{total time} - \text{scheduled time})] * 100$) [AVL].

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

Increasing the overall performance delivered to users has been an on-going effort since the introduction of service level management. Availability/Reliability averaged per quarter across the whole infrastructure have been both steadily increasing from 2008 by approximately 1% per year, moving from 91.9%/93.3% during May 2009 – April 2010 (last year of EGEE-III), to 94.50%/95.42% during May 2011 – March 2012 (second year of EGI-InSPIRE).

Performance exceed the PY2 Reliability target (91%) – see Table 9. Even if performance of Resource Centres is gradually increasing as illustrated in Figure 16, availability and reliability still reflect the on-going increasing maturity of the overall infrastructure and can be adversely affected by changes (e.g. upgrades in software, new Operations Centres coming into business, new production sites being certified).

Several NGIs already integrated during PY1 were affected by periodic performance instability because of lack of expertise in technical services being deployed. A support action was kicked off during PY2 in collaboration with the Greek JRU to technically support those NGIs in the South East Europe region. In addition, a training event for site administrators is being organized in collaboration with the technology providers. Other RCs frequently affected by very low performance were in the Asia Pacific region. Asia-Pacific is currently the largest federated operations centre encompassing 10 different countries, internally showing very different maturity levels in grid knowhow.

The trend of the overall EGI availability and reliability is plotted in Figure 16, which plots the average quarterly availability and reliability of RCs from May 2010. For example, the deep recorded from August to November 2011 reflects the transition of various large federated Operations Centres evolving towards a set of operationally independent NGIs. EGI currently comprehends 347 certified Resource Centres. These increased by 9.5% since the end of EGEE-III. Given this expansion trend, the definition of managed process to gradually and safely bring new RCs from test to production, is now part of the activity roadmap of the EGI operations community.

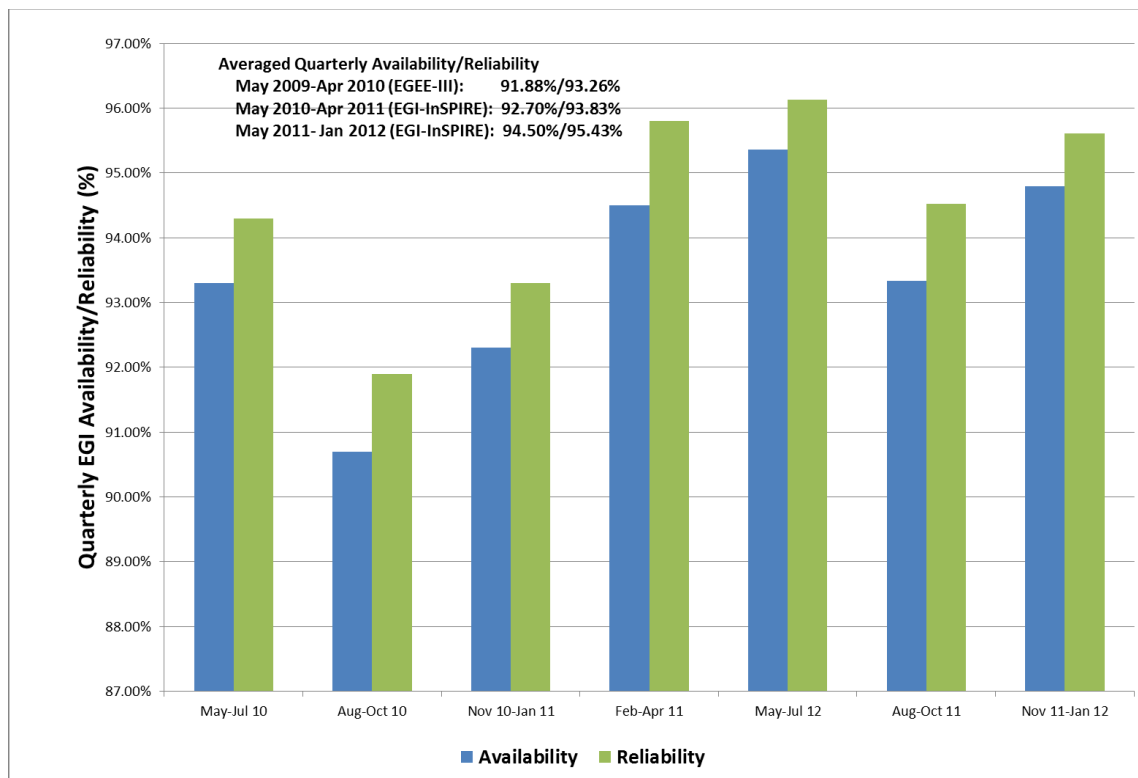


Figure 16. Quarterly availability and reliability of resource centres averaged across EGI from May 2010 to end of PQ7.

A new profile used for measurement of RC availability and reliability was adopted started on 01 January 2012 after an impact assessment conducted for two months at the end of 2011. The new profile name is ROC_CRITICAL¹⁴ and it replaced the previous profile WLCG_CREAM_LCGCE_CRITICAL¹⁵. Reasons for updating the profile were the need to include more monitoring probes (e.g. for site-BDII services), and to decouple of the OPS VO profile used by EGI from the one used by WLCG.

6.3 RP Performance

User experience not only depends on resource-access services, 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 to access distributed resources.

RP performance is reported monthly. Purpose of this reporting is to check the availability and reliability of core services operated by NGIs and EIROs, which is typically highly critical as these services provide access to RC services, and are often shared across multiple user communities. In

¹⁴ http://grid-monitoring.cern.ch/myegi/sam-pi/metrics_in_profiles?vo_name=ops&profile_name=ROC_CRITICAL

¹⁵ http://grid-monitoring.cern.ch/myegi/sam-pi/metrics_in_profiles?vo_name=ops&profile_name=WLCG_CREAM_LCGCE_CRITICAL

order to enhance robustness and performance, these services are frequently constituted of distributed physical instances deployed across multiple RCs. In this case, performance results from the compounded availability of the service physical instances.

6.3.1 Availability and Reliability

Current availability and reliability reports include statistics for the information discovery services (top-BDII). The set of monitored core services will be gradually extended to include workload management systems, file catalogues, VO management services etc.

As shown in Figure 17, currently 77% of the Resource infrastructure Providers operate top-BDII services with an average monthly availability that exceeds 97%. France, Greece, Italy, The Netherlands and Turkey are the NGIs whose top-BDII delivered 100% availability during the while reference period.

Starting in 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 is to further enhance the performance offered to end-users.

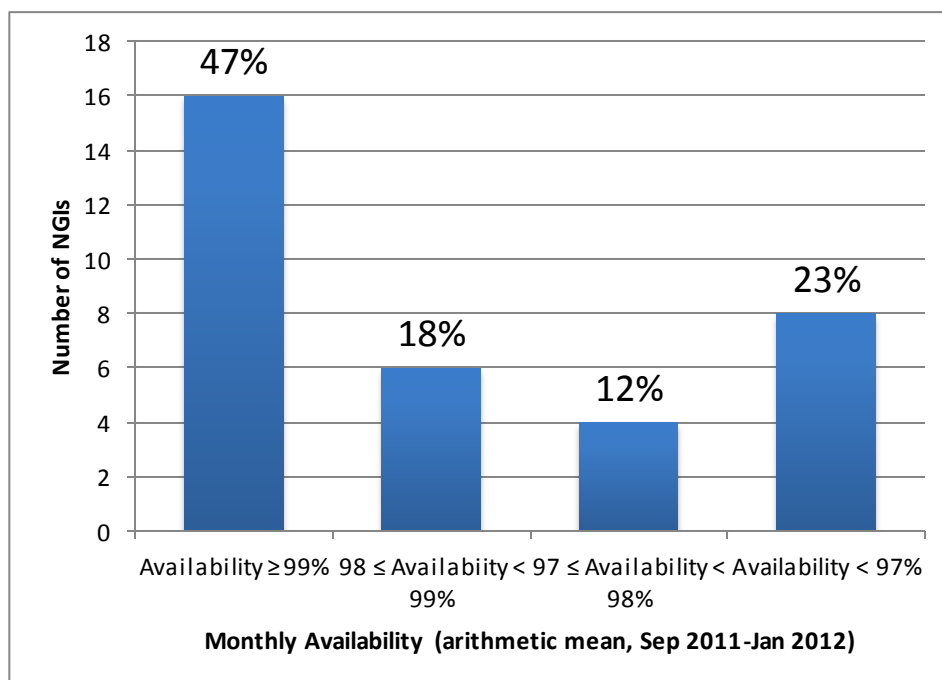


Figure 17. Classification of NGIs by the monthly availability provided by their top-BDII services. Here NGIs are classified according to the average monthly availability calculated during the reference period September 2011 – January 2012 (end of PQ7).

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 manual [MAN05].
- The list of authoritative top-BDII was collected and their configuration was assessed.

- The list RCs making use of the CERN top-BDII as primary instance was collected and the NGIs were requested to support the administrators change configurations, so that the correct authoritative instance is used instead.

6.3.2 ROD Performance Index

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

The ROD performance index¹⁶ is the sum of the number of ticket expired in the operations dashboard daily, and the number of alarms older than 72h appearing in the operations dashboard daily

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

The ROD performance index was introduced to complement the ROD workload metric, which is highly dependent on the size of the operated infrastructure.

The chart in Figure 18 shows the NGIs' Regional Operator on Duty (ROD) average monthly workload, which is the sum of all the items the ROD team has to deal with in a month. These items can be either alarms in the Operations Dashboard older than 24h or any ROD ticket generated through the operations dashboard.

The chart in Figure 19 shows the percentage of Operational Dashboard alarms closed to 'OK' status. The values are the arithmetic means calculated over quarters PQ5-PQ6-PQ7. Alarms should be closed only when the problem is solved. There are only a few cases where an alarm can be closed in 'non OK' status, (e.g. second test run is successful, but the result is not correctly propagated to the dashboard, therefore the dashboard shows the alarm in a 'error' status, while it is actually in an 'ok' status). A ROD team scores a high percentage in this metrics when active support is provided, and alarms are properly handled and closed only when the issue is solved.

The average quality performance of the ROD teams was consistently above 90% in the last three quarters.

¹⁶ https://wiki.egi.eu/wiki/Grid_operations_oversight/ROD_performance_index#Definition




Figure 18. ROD average monthly workload (from PQ5 to PQ7).




Figure 19. ROD quality metric (monthly average from PQ5 to PQ7).

6.4 NGI Support Actions

As shown in the previous sections, several NGIs suffered from low availability. A support action was delivered in collaboration with the Greek JRU to help the identified NGIs to solve the technical issues faced in running RP core services (including the NGI SAM service) and the RC services.

- Albania: this NGI is still not contributing resources to EGI. The Greek JRU was requested to provide support in bootstrapping the operations centre and one RC, however, the issue was then identified to be related to lack of local staff. This problem is still open.
- Armenia: the NGI experienced problems with the configuration of the NGI SAM service, which were fixed. Errors in publishing installed capacity were also corrected, as well as basic configuration problems in the compute elements and accounting clients. At the time of writing NGI_ARMGRID comprehends three production sites, whose performance stabilized in March, and one production site still affected by extremely low performance, which is now back to uncertified status.
- Bulgaria: an increasing large fraction of sites has been periodically affected by sporadic low performance. The reason for this is being investigated.
- FYR of Macedonia: the infrastructure currently comprehends three sites. Two of these were severely underperforming during PQ7 and PQ8. Problems were reported to be due to staff turnover and changes in the hosting infrastructure, new resources are being procured and integrated.
- Moldova: the first production site was certified in March 2012. The operations manager of Moldova was funded by EGI.eu to participate to the Community Forum 2012.
- Montenegro was affected by low availability during PQ7 due to top-BDII hardware problems. The service was replaced and since then, performance has stabilized during PQ8.
- The operations manager of Romania was funded by EGI.eu to participate to the Community Forum 2012.

The tools and procedures in place that allow a RC to safely move from testing to production status are being reviewed. These are currently not sufficiently flexible, and experience in running a grid infrastructure at a RC level is being currently gathered while running a certified production site. A training programme for site administrators is being planned in collaboration with EMI,

7 VO SERVICES

The EGI VO Services [VOS] – defined at the beginning of EGI-InSPIRE as part of NA3 and now migrated to SA1 – aim at supporting VOs in the whole process of start-up, management and operation, pointing out to tools, services, documentation and guidelines to maximize the usage of the resources, easing service deployment, and bridging the VO community with the infrastructure need.

VO services are mature enough to be supported by NGI operational teams and the expertise on operating those services is also widely available in the operations community. NGIs are responsible for supporting and operating the developed VO services, and of contributing to the maintenance of the existing documentation base.

EGI operations are responsible for:

- The operation of VO middleware services and operational tools including VO-tailored SAM instances, which support VO-dedicated services and VO-centric testing and monitoring [VSAM]. VO SAM instances are distributed and typically operated either by VO partners or RPs.
- The operation of a VO administration dashboard¹⁷ for the configurable integration of EGI and VO tools [VOA]. The VO administration dashboard is technically provided by LIP.
- LFCBrowseSE [LFCB] support, a tool to improve file catalogue views.
- Testing of the VO Operation Dashboard [VOD], a dashboard that based on VO SAM results can raise alarms that VO Operations team can assess.
- Contributing and enhancing documentation [DOC] regarding VO activities and tools.

7.1 Status

VO SAM: is now a flavour officially supported by the SAM product team within EGI-InSPIRE JRA1, and it will continue to be enhanced with the introduction of a new component: the Profile Management Database (POEM) substituting the Metric Description Database (MDD), allowing higher flexibility to describe existing metrics¹⁸ and group them in order to run dedicated tests.

VO Admin Dashboard is a highly configurable integrator dashboard for project and VO tools, in production since mid-May 2011. It is a tool tailored for VO management activities, supported and operated by the IBERGRID federation at LIP, and open for use to all interested communities. New tool enhancements are only provided on a best effort basis.

LFCBrowseSE is a tool to improve file catalogue views, it is currently in its third release. It is supported on a best effort basis by the IBERGRID federation at UPV, and it extensively used by the LSGC VRC for SE decommissioning activities.

VO Operational Portal is a module fully integrated with the main Operations Portal, foreseen to empower VO operators with a service allowing the collection of alarms in case of critical failures of user workflows, and thus contributing to the enhancement of the availability of the infrastructure perceived by end-users. It has been developed by JRA1 following requests from the interested user

¹⁷ <https://vodashboard.lip.pt/>

¹⁸ https://wiki.egi.eu/wiki/SAM_Tests



communities and according to an approved work plan [WPL] defined in collaboration with JRA1 and the user communities. It is presently in prototype phase, it is being tested.

7.2 Utilization

VO SAM is being used by a number of different communities, including phys.vo.ibergrid.eu (one of the top 10 VOS operating in EGI infrastructure since January 2012), life.vo.ibergrid.eu, ict.vo.ibergrid.eu and fusion. Although not registered in GOCDB, other communities are also using similar services, such as We-NMR and HealthGrid (on behalf of the Life-Science Grid Community VRC) for which consultancy has been provided. SA1 plans to launch a campaign suggesting the registration of the services in GOCDB so that one could assess the relevance and importance of the service inside user communities.

VO Admin Dashboard: The tool is operated since mid-May 2011. According to Google Analytics from 1st May 2011 to the end of March 2012 the tool got 237 unique visitors, 1056 visits and 4601 page view. Given the scope for which the tool was developed (it is for VO Managers and Administrators), we consider this result as very good, meaning that the tool is being access by VO users as well. Moreover, the fact that the average visit duration is of more than 4 h and that 77.65% represents visitors that have returned, means that the tool is fulfilling its purpose, and it is really used on the VO daily work.

LFCBrowseSE was developed following a request from Life Science VRC, and it extensively used by this community for SE decommission purposes. It has a valuable potential to be used by other communities.

The VO Operations Portal is currently under test but already supports 4 Vos: biomed, phys.vo.ibergrid.eu, prod.vo.eu-eela.eu and vo.plgrid.pl.

8 GRID SERVICES

In this section we review the status of deployment of different software platforms across EGI. As indicated in Table 10, the set of software platforms that are being successfully integrated with EGI operations, currently encompasses EMI software (ARC, dCache, gLite, UNICORE), GLOBUS being maintained, released and supported by the IGE project, QoSGrid supported by PL-Grid¹⁹, and Desktop Grid software released and supported by the EDGI project²⁰.

The list of production end-point services and their type can be obtained programmatically from the GOCDB programmatic interface²¹.

8.1 Integrated Software Platforms

Table 10. Deployment of integrated software platforms across EGI

Integrated Grid Platform	Number of countries	Countries
ARC	8	Denmark, Estonia, Finland, Latvia, Norway, Slovenia, Sweden, Switzerland. During 2012 the ARC deployment is expected to be extended to the Ukrainian NGL.
Desktop Grid (experimental phase)	1	Hungary
GLOBUS	3	GridFtp: Finland, Germany, Spain GRAM: Germany During 2012 the GLOBUS production deployment is expected to be extended to The Netherlands and United Kingdom.
QosCosGrid (experimental phase)	1	Poland
UNICORE	2	Germany, Poland

Currently the grid middleware stacks fully deployed in the infrastructure are: gLite, ARC, GLOBUS and UNICORE. In the last year the integration level of the various stacks was consolidated, even though it cannot be considered totally complete yet. Currently the EGI service registry (GOCDB) defines the service types necessary to register services from all the stacks.

¹⁹ <http://www.egi.eu/community/collaborations/MAPPER.html>

²⁰ <http://www.egi.eu/community/collaborations/EDGI.html>

²¹ <https://wiki.egi.eu/wiki/GOCDB/services>

Nagios probes are available to test the service status of the aforementioned stacks. For UNICORE and GLOBUS the full integration in the SAM infrastructure and in the operational shifts of the operations centres, will be completed in the coming months.

Accounting integration is still in progress for UNICORE and Globus, while ARC and gLite computing resources have been accounted for their usage from the beginning of EGI-InSPIRE. The Accounting Task Force of the TCB²² is responsible of leading the extension of the current EGI accounting infrastructure to encompass peer grids and new integrated infrastructures.

gLite is largely the most deployed middleware, but the number of operations centres supporting non-gLite stacks slightly increased during last year. As shown in Figure 20, 4 operations centres are deploying ARC middleware, namely: NGI_NDGF (including Denmark, Estonia, part of the Finnish resources, Latvia, Norway, Sweden), NGI_CH, NGI_FI, NGI_SI. UNICORE is supported by two operations centres NGI_DE and NGI_PL, and Globus middleware is deployed by NGI_DE.




Figure 20, deployment of the four reference grid middleware stacks across the EGI-InSPIRE operations centres, March 2012 (source GOCDB)

F. For EDGI middleware, currently already supported in GOCDB, Nagios probes will be available in SAM Update 18. QCG/MAPPER software, already supported in GOCDB, and cloud middleware (e.g. StratusLab or Open Nebula) in the EGI operational fabric are being experimented to fulfil the needs of new user communities.

²² https://wiki.egi.eu/wiki/TCB:Accounting_Task_Force




Figure 21, Number of instances of the different implementations of the compute capability, across the EGI_InSPIRE partners and the integrated Resource infrastructure Providers, March 2012 (source: GOCDB)

Figure 21 shows the distribution of the deployed computing element instances among the different alternative software stack implementations.

Lcg-CE reached end of support at the end of April 2012. Lcg-CE instances have considerably reduced in number in the last year, as resource centres are migrating to other supported CE implementations. However, the phasing out is still in progress. Two resource centres are deploying the Globus Gram, and there are 31 ARC-CE instances deployed in the infrastructure.

8.2 Deployment of gLite releases

The majority of the gLite 3.1 products are no more supported²³ and Scientific Linux 4 – on which gLite3.1 components must be deployed – ended its security support in February 2012. In the February 2012 meeting the Operations Management Board approved the following decommissioning plan: by 30/09/2012 all the instances of the following products need to be phased out:

- lcg-ce
- any site-BDII instance earlier than 3.2.10-1 (i.e. not GLUE 2.0 compatible)
- all gLite 3.1 currently unsupported products – including all gLite 3.1 products with the exception of gLite-Cluster, glite-LSF_utils, glite-SGE_utils, glite-TORQUE_utils.

After 30/09/2012 sites still hosting one of the services in the above list will be suspended.

gLite 3.2 was supported until end of April 2012, consequently many sites will likely migrate from gLite3.1 directly to the equivalent EMI-1 components (supported until April 2013) or to the upcoming EMI-2 that will support SL5 and SL6.

Figure 22 shows the deployed instances, grouped by the gLite and EMI major releases, of the main compute and storage management gLite components: Worker Node, Workload Management System, DPM and LFC.

²³ gLite 3.1 support calendar: <http://glite.cern.ch/R3.1/>


Most of the worker nodes are gLite 3.2, only 7% of the sites are deploying worker nodes released in gLite 3.1. As to the number of cores deployed in those sites, the percentage of computing capacity delivered through gLite 3.1 WNs is around 1% of the total. The figures for worker nodes may not be accurate, as for technical limitations it is not possible to gather complete information about RCs deploying different releases of worker nodes.

CREAM (372 total deployed instances) has also a majority of gLite 3.2 instances (78.8%) and EMI (17.7%). Only a small fraction is still from a gLite3.1 release (3.5%).

EMI WMS release (35.3%) has a good fraction of the total instances deployed (total 150), even though most of the instances (64.7%) are still gLite3.1 (WMS has not been released in gLite3.2). EMI WMS was released for the first time in UMD with Update 1.5.

There are no EMI instances for central LFCs. 63% deployed instances are gLite3.2 and the remaining 37% are gLite3.1. 41 instances in total are present in the production infrastructure.

DPM seems to be the service with the highest percentage of EMI release installations, unfortunately the same service versions were released in both gLite3.2 and EMI (same problem with gLite 3.1 and gLite3.2), therefore, getting precise figures from the information discovery system according to the publish number of version it is not possible. The DPM instances are 206 in total, of which 43.3% are EMI , 58.3% gLite 3.2 and 1.2% gLite 3.1.






Figure 22. Deployment of the different (gLite and EMI) releases for the services: WN, CREAM, WMS, LFC and DPM, March 2012 (source: information discovery system).

8.3 Core Middleware Services

Core grid 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: LFC (file catalogue), WMS (workload management), Top-BDII (information system top-level cache) and VOMS (VO membership, attribute management).


Figure 23 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 are not available in the services registry (GOCDB).

In March 2012 EGI integrated infrastructure comprehends 359 core services: 64 VOMS instances, 150 WMS, 41 LFC and 104 Top-BDII. The top five NGIs/EIROs as number of core services instances deployed are: Italy (45 instances), CERN (43 instances), France and Germany (29 instances each) and Spain (27 instances). 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 these services. There is still a large number of gLite 3.1 instances, since WMS was not released in gLite 3.2.

Top-BDII is deployed by all the NGIs who are deploying gLite middleware, since it is a critical component for the gLite workflow. Currently 31 NGIs are deploying at least one Top-BDII. Given its criticality, starting from September 2011 Top-BDIIs availability is monitored, and starting from January 2012 NGIs are requested to meet the target of 99% monthly availability. Currently 14 NGIs and EIROs are deploying Top-BDII with a high availability configuration.


The VOMS chart has no information about the deployed versions, because some instances do not publish reliably themselves in the information system. To get accurate data about distribution of the VOMS services, the instances list was fetched from GOCDB.



(b)
Workload
Management
(WMS)



(c)
VO
membership
(VOMS)
Source:
GOCDB






Figure 23, Number of core services instances deployed within the EGI-InSPIRE integrated infrastructure. Where not specified, source: information system (March 2012)

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 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.

The Staged Rollout workflow introduced during PY1, was refined during the first year of EGI-InSPIRE, this has been done in parallel with the construction of the Staged Rollout infrastructure, which is being gradually expanding reflecting the deployment needs of VRCs and NGIs.

Table 11. Overview of EGI-InSPIRE Staged Rollout metrics.

Metric	PQ5	PQ6	PQ7	PQ8
	UMD 1.0.0	UMD 1.1.0/1.2.0/1.3.0	UMD 1.4.0/1.5.0	
Total number of components tested/rejected	54/2	30/3	30/3	TBD
Number of staged rollout tests undertaken	81	49	50	TBD
Number of EA teams	46	52	56	TBD

As shown in Table 11 the largest number of products was tested in PQ5 in preparation to the release of UMD 1.0 (81 tests in total resulting in 2 products rejected). This number was gradually reduced in the following quarters, as subsequent UMD updates only included a subset of products being updated. The number of participating Early Adopters has been progressively increasing to test a growing set of products from EMI, IGE and EGI-InSPIRE JRA1 (operational tools), and it currently amount to 56 teams.

The number of tests performed from PQ1 to PQ2 by NGIs and EIROs is plotted in Figure 24. Participation of partners to Staged Rollout activities will be reviewed in PY3 in order to ensure enough coverage of testing of EMI components for different platforms (sl5, sl6 and Debian) and IGE. In addition, the current staged rollout process will be further enhanced.




Figure 24. Number of Staged Rollout tests performed from PQ1 to PQ8 by NGIs/EIROs.

10 CONCLUSIONS AND FUTURE WORK

The Resource Infrastructure status satisfactorily met the infrastructure targets defined for the first year of EGI-InSPIRE (number of RCs integrated with -1.21% deviation, number of job slots offered, EGI monthly reliability, usage, MPI support, number of HUC and non-HUC VOs and HPC resources). The Desktop Grid integration is being piloted in Hungary and is currently under finalization.

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

This objective was successfully met by completing the transition process started in PY1 that evolved the EGEE federated Operations Centre into independent NGIs. Two new Resource infrastructure Provider MoUs were signed with the South African Grid Initiative and the Ukrainian National Grid; expansion in these regions is expected in PY3. Moldova became operational in March 2012, and a grid infrastructure is emerging in Azerbaijan. In PQ2 the number of production RCs increased to 352. The installed capacity and Resource Centres grew considerably to comprise 270,800 logical cores (+30.7% yearly increase), 2.96 Million HEP-SPEC 06 (+49.5%), 139 PB of disk space (+31.4%) and 134.3 PB of tape (+50%). During PY3 expansion will be mainly driven by the integration of new providers and will depend on the capability of integrating new software platforms. The average EGI Availability and Reliability from May 2011 to March 2012 are 94.50% and 95.42% respectively. PY2 yearly EGI averaged Availability increased by +1.8%, and Reliability by 1.5%. A support action was launched in PQ6 in collaboration with the Greek JRU to technically support several infrastructures requiring consolidation.

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

In PY2 the responsibility of providing VO services was migrated to the EGI.eu operations team and the NGIs. VO support includes existing SA1 VO services provided by NGIs including support through the EGI helpdesk, the operation of software platforms dedicated to VOs (VO Management Services, user identity provisioning, VO grid services etc.), and the operation of tools to assist VO administration and monitoring.

The overall number of international and national VOs registered in the Operations Portal at the end of March 2012 amounts to 226 (3.20% from March 2011), including 20883 registered users (14.30% increase from March 2011).

Overall resource utilization has been satisfactorily progressing confirming the trends of PY1. The yearly increase of the total number of jobs executed in the infrastructure in the period May 2011-April 2012 amounts to +46.42% of the yearly job workload done from May 2010 to April 2011. The PY2 overall quantity of EGI computing resources used amounts to 10.5 Billion HEP-SPEC06 Hours. The PY2 workload was generated by 492.5 Million jobs (1.35 Million Job/day on average). The High-Energy Physics discipline (contributing 39.25% of the user community) is still expanding in resource utilization, and the used normalized CPU wall time increased from 91.13% (April 2011) to 93.60% (April 2012) of the overall amount of EGI resources used. In particular, the used normalized CPU wall time yearly increased by +48.82%, while the job rate yearly increased by +57.06%. Astronomy Astrophysics and Astro-particle Physics are the second community in terms of used normalized CPU wall clock time, which now amounts to 2.25% of the overall EGI used CPU wall clock time, showing a

+117.79% yearly increase from April 2011. Life Sciences are the third community for usage (1.30% of the overall EGI used normalized CPU time).

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

Currently the grid middleware stacks fully deployed in the infrastructure are: gLite, ARC, GLOBUS and UNICORE. The integration of other software platforms is in progress (Desktop Grids and QCG/MAPPER software). The availability of Desktop Grid resources will allow the integration of additional capacity for opportunistic usage of the infrastructure. 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.

- Objective 5 (O5): *Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure so as to provide transparent access to all authorised users.*

The “Resource Infrastructure Provider Operational Service Agreement” [RPO] was approved in October 2011 to facilitate the exchange of operational services and the integration between the EGI-InSPIRE infrastructure and those operated by internal and external partners. A revised version was subsequently approved in March 2012. In addition, two Resource infrastructure MoUs were approved with the South African Grid Initiative and the Ukrainian NGI. The successful integration of new infrastructures is a pre-requisite for the support of international user communities and to foster collaboration between scientists across the world.

- 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.*

The support of parallel jobs has been consolidating and expanding across the infrastructure. The full integration of UNICORE resources into EGI – thanks to the harmonization of Grid middleware and operational interfaces – will further foster the expansion of the high-performance capabilities of EGI. In addition, thanks to the collaboration with the MAPPER project, EGI operations are piloting the support of applications requiring a loosely and/or coupled usage of HTC resources from EGI and HTP resources from PRACE. Various are the disciplines that can benefit from this approach, including earth science, meteorology and seismology. The MAPPER integration use cases are being used to define the EGI/PRACE common operations roadmap. A PRACE workshop was held in Munich during the Community Forum 2012²⁴ to discuss this, in order to understand what other user communities would benefit from coupled usage HPC and HTC resources. The definition of the EGI/PRACE collaboration framework is in progress and various services are being investigate.

²⁴ PRACE workshop agenda at the Community Forum 2012:

<https://www.egi.eu/indico/sessionDisplay.py?sessionId=28&confId=679#20120328>


11 REFERENCES

APL	Roger, W., Jiang, M., Del Cano, C. et alt.; <i>An APEL Tool Based CPU Usage Accounting Infrastructure for Large Scale Computing Grids</i> , in the Proc. of Int. Symposium on Grid Computing 2010, Taipei (TW), Mar 2010 (http://event.twgrid.org/isgc2010/abstract/AnAPELToolBasedCPUUsageAccountingInfrastructureforLargeScaleComputingGrids.pdf)
AVL	EGI availability and reliability (https://www.egi.eu/infrastructure/Figures_and_utilisation/Availability_reliability.html)
D4.3	EGI Operations Architecture: Grid Service Management Best Practices, EGI-InSPIRE Deliverable D4.3, Dec 2011 (https://documents.egi.eu/document/763)
D4.4	Security Risk Assessment of the EGI Infrastructure, EGI-InSPIRE Deliverable D4.4, Nov 2011 (https://documents.egi.eu/document/863)
D7.2	Annual Report on Operational Tool Maintenance, EGI-InSPIRE Deliverable D7.2, March 2012 (https://documents.egi.eu/document/1063)
DOC	EGI Documentation of VO Services (https://wiki.egi.eu/wiki/VO_Services/Documentation)
GLO	EGI Glossary (https://wiki.egi.eu/wiki/Glossary)
GTR	EGI Global Task Review, EGI-InSPIRE Milestone MS115, March 2012 (https://documents.egi.eu/document/961)
HS06	https://wiki.egi.eu/wiki/HEP_SPEC06
LFCB	LFCBrowseSE (http://wiki.healthgrid.org/Biomed-Shifts:LFCBrowseSE)
MAN05	Top-BDII High Availability, EGI Manual MAN05 (https://wiki.egi.eu/wiki/MAN05)
MAP	Multiscale Applications on European e-Infrastructures (http://www.mapper-project.eu/)
MS115	EGI Global Task Review, EGI-InSPIRE Milestone MS115, March 2012 (https://documents.egi.eu/document/961)
MS116	EGI NGI International Task Review, EGI-InSPIRE Milestone MS116, March 2012 (https://documents.egi.eu/document/963)
MS414	Integrating Resources into the EGI Production Infrastructure, EGI-InSPIRE Milestone MS414, Sep 2011 (https://documents.egi.eu/document/650)
MS418	Operational Level Agreements (OLAs) within the EGI production infrastructure, EGI-InSPIRE Milestone MS418, April 2012 (https://documents.egi.eu/document/1057)
PERF	EGI service level monthly reports (https://wiki.egi.eu/wiki/Performance)


PRO	EGI monitoring profiles (https://wiki.egi.eu/wiki/SAM#Profiles)
QR7	Quarterly Report 7, EGI-InSPIRE Milestone MS114 (https://documents.egi.eu/document/999)
RCO	Resource Centre Operational Level Agreement v1.1, March 2011 (https://documents.egi.eu/document/31)
RPO	Resource infrastructure Provider V1.1 (https://documents.egi.eu/document/463)
SAG	South African Grid Initiative – SAGrid (https://documents.egi.eu/document/495)
SAM	Service Availability Monitoring (https://wiki.egi.eu/wiki/SAM)
SAMV	Service Availability Monitoring for VOs (https://wiki.egi.eu/wiki/Services_and_Tools_Portfolio)
SRW	Staged Rollout: https://wiki.egi.eu/wiki/Staged-Rollout
STR	EGI Strategic Plan, EGI-InSPIRE Deliverable D2.30, April 2012 (https://documents.egi.eu/document/1098)
UMD	Unified Middleware Distribution (http://repository.egi.eu/category/umd_releases/distribution/umd_1/)
VOA	VO Administration Dashboard: https://wiki.egi.eu/wiki/VO_Services/VO_Admin_Dashboard
VOD	VO Operations Dashboard (https://operations-portal.egi.eu/voDashboard)
VOS	EGI VO Services (https://wiki.egi.eu/wiki/VO_Services)
VSAM	VO SAM: https://wiki.egi.eu/wiki/VO_Services/VO_Service_Availability_Monitoring
UNG	MoU between EGI.eu and BCC – Ukraine (https://documents.egi.eu/document/8560)
WPL	VO Operations Portal work plan, March 2012 (https://documents.egi.eu/document/1055)

12 ANNEX I. VO DISTRIBUTION PER DISCIPLINE


12.1 Astronomy Astrophysics and Astro-particle Physics




12.2 Computer Science and Mathematics




12.3 Earth Sciences



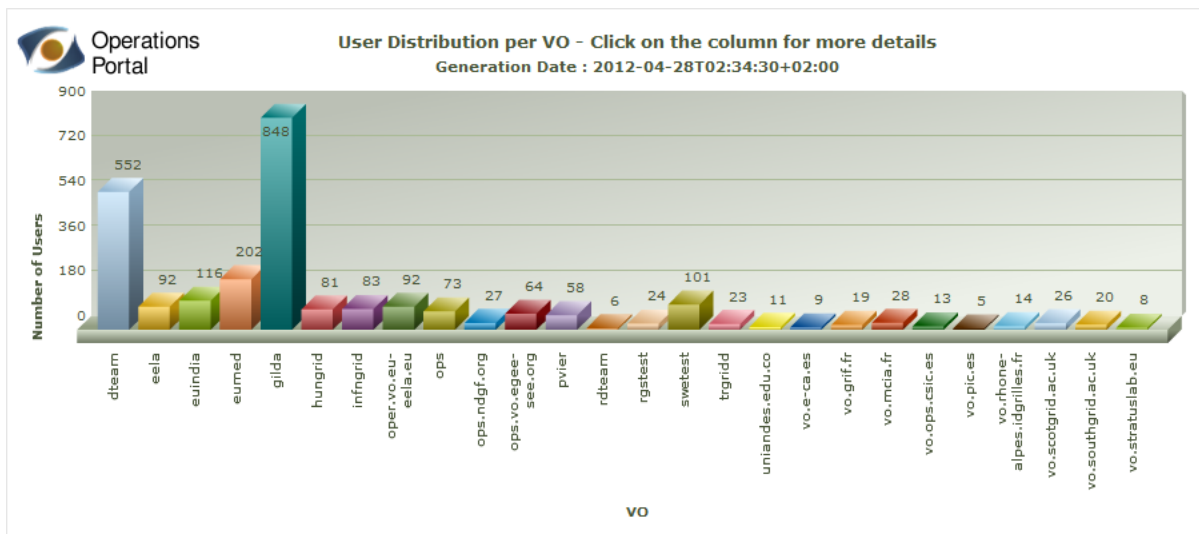
12.4 Fusion




12.5 High Energy Physics




12.6 Infrastructure



12.7 Life Sciences



12.8 Multidisciplinary VOs



12.9 Other Disciplines

