EGI-InSPIRE

Annual Report on the EGI Production Infrastructure

EU DELIVERABLE: D4.2

Document identifier:	EGI-D4.2-v1.3
Date:	**16/05/2011**
Activity:	**SA1**
Lead Partner:	**EGI.eu**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/218

Abstract

This deliverable provides information about the transition and status of the EGI-InSPIRE Resource Infrastructure as of the end of the EGEE-III project to March 2011. The document describes the operational structures and installed compute and storage capacity operated by EGI-InSPIRE partners and by a set of integrated external providers. Utilization, performance and the deployed are presented.

I. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Tiziana Ferrari	EGI.eu/SA1	02/04/2011
Reviewed by	**Moderator:** G. Sipos **Reviewers:** T. Glatard, M. Chatziangelou	EGI.eu/NA3 CNRS/NA3 IASA/NA3	06/05/2011
Approved by	**AMB & PMB**		14/05/2011

II. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	02 April 2011	Draft for internal/external review v 1.0	T. Ferrari and P. Solagna/EGI.eu
2	20 April 2011	Draft v.1.1	T. Ferrari/EGI.eu
3	04 May 2011	Draft v.1.2	T. Ferrari, P. Solagna/EGI.eu
5	12 May 2011	Final	T. Ferrari, P. Solagna/EGI.eu

III. APPLICATION AREA

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

IV. DOCUMENT AMENDMENT PROCEDURE

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

V. TERMINOLOGY

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

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 supports 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 is also 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 collects user requirements and provide support for the current and potential new user communities, for example within the ESFRI projects. Additional support is also 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 Resource 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.

VI. EXECUTIVE SUMMARY

The first year of EGI-InSPIRE was dominated by a structural change of the operational organization of the Resource Infrastructure. The 14 EGEE Regional Operations Centres that were active in April 2010 have developed into a much larger group of smaller Operations Centres (31 in total). These Operations Centres typically serve a single country (22 single NGI Operations Centres, 5 federated European NGI Operations Centres, and four non-European ones). At the end of March 2011 EGI comprised one European Intergovernmental Research Organization - EIRO (CERN) and 40 National Grid Infrastructures of which two (Albania and Moldova) are expected to join in 2011. This transition did not significantly affect the overall average EGI monthly performance and the average monthly reliability met the project target (90 % reliability).

The Resource Infrastructure has been steadily expanding well beyond the official project targets in terms of Resource Centres (300) and compute resources installed (200,000 cores), and has been gradually enhancing its level of integration with multiple middleware stacks and its support to parallel jobs. Utilization – mainly driven by the High-Energy Physics discipline – has been dramatically increasing. New processes have been put in place for validating new software releases before large-scale deployment, and the so-called Staged Rollout infrastructure is consolidating and expanding.

In March 2011 the compute resources contributed by the EGI-InSPIRE partners amounted to 207,203 logical CPUs (1.98 Million HEPSEPC 06), this corresponds to a +7.9 % increase since April 2010. The installed capacity increased to from 207,203 to 308,583 logical CPUs if we also consider the integrated Resource Infrastructures (e.g. Canada and Latin America), and peer Grids (Open Science Grid[1] – OSG – and South Africa Grid - SAGrid). The estimated storage capacity amounts to 101 PB of disk capacity and 80 PB of tape capacity.

The level of penetration of support to parallel application has been increasing. In particular, the number of integrated high-performance clusters amounts to 54 units at the end of Project Quarter (PQ) 3. In addition, Message Passing Interface jobs are supported by 90 Resource Centre, which corresponds to 26.47 % of the total number of Resource Centres at the end of PQ3.

The number of Resource Centres has been increasing (this is a constant trend since May 2008), and now they amount to 332 Resource Centres distributed among 58 countries and one EIRO. Out of these, resources are directly contributed by EGI-InSPIRE partners in 45 countries and one European Intergovernmental Research Organization (CERN). The remaining 13 countries contribute resources through Resource Infrastructures that are operated outside of EGI-InSPIRE. These resources are fully integrated with the EGI Services Infrastructure.

The monthly Availability and Reliability have been fluctuating in the range [84.4, 95.85] %, without a specific trend. The average monthly EGI Availability and Reliability are 90.7 % and 91.9 % respectively. Both values are in-line with the project target Reliability of 90 %. EGI performance was mainly affected by the decommissioning of legacy federated Operations Centres[2], which generated a large amount of new centres. Since the start of EGI-InSPIRE only six Resource Centres in total were suspended because of performance issues or unresponsiveness to trouble tickets.

[1] http://www.opensciencegrid.org/

[2] Operations Centre decommissioning procedure: https://wiki.egi.eu/wiki/PROC03

According to the Operations Portal, at the end of March 2011 the EGI users were 13319 (9.5% increase from March 2010), while the VOs were 186 (+17.4% increase from March 2010). Among all users the High-Energy Physics (HEP) community have greatly increased the generated workload. From May 2010 to March 2011 the average number of jobs executed per month amounted to 25.70 Mjob/month including all VOs (1.52 Mjob/month for non-HEP VOs), while the average number of executed job per day greatly exceeds the target of 500,000 units and it amounts to 933,000 jobs/day (55,200 jobs/day for non-HEP VOs). If we consider compute utilization across all Resource Infrastructures (from EGI-InSPIRE partners and integrated providers), then the utilization of the integrated EGI-InSPIRE infrastructure (including integrated Resource Infrastructure Providers that are consumers of EGI services) is 70.8%, while the OSG utilization amounts to 29.2%.

The Staged Rollout Infrastructure currently comprises 39 Resource Centres, and several Early Adopters are participating to the Staged Rollout of two or more software components. Presently the Staged Rollout Infrastructure is complete for the gLite distribution (3.1 and 3.2), the ARC middleware stack, most of UNICORE and Globus components, and one Operational Tool (the SAM framework and Nagios probes). From the beginning to the end of PQ3, 32 patches were tested by one or more Early Adopters before being deployed by all production Resource Centres.

TABLE OF CONTENTS

1 INTRODUCTION

The European Grid Infrastructure (EGI) provides a sustainable foundation to the evolution and operation of a federated distributed computing infrastructure to support the European Research Area.

The production infrastructure rests on two pillars: the *Resource Infrastructure* and the *Service Infrastructure* [D4.1]. The Resource Infrastructure has a tiered architecture. Firstly, resources are geographically distributed, and are contributed by *Resource Centres*. A Resource Centre is the smallest resource administration domain within EGI. A *Resource Infrastructure* federates one or more Resource Centres to constitute a homogeneous operation domain, and the federation of Resource Infrastructures of EGI constitutes the EGI Resource Infrastructure. On the other hand, the operational Service Infrastructure enables a secure, interoperable and reliable access to the Resource Infrastructure. EGI operational services are provided locally by *Operations Centres* and globally by *EGI.eu*. Local and Global Services are complementary.

This deliverable focuses on the status of the Resource Infrastructure covering the period from May 2010 to March 2011. The Service Infrastructure is not described in this document. The interested reader is encouraged to get information about the EGI Global and Local Services in [MS108], [MS109].

Resources are geographically distributed, and are contributed by *Resource Centres* (Figure 1). A Resource Centre is the smallest resource administration domain within EGI. A Resource Infrastructure federates one or more Resource Centres to constitute a homogeneous operations domain. The Resource Infrastructure usually encompasses heterogeneous resource types. Currently these are mainly high throughput computing, high performance computing, and storage, which are seamlessly made accessible through the deployment of standard interfaces and gateways provided by various Grid middleware stacks such as ARC, gLite, UNICORE and Globus. New resource types will be integrated as technologies mature during the project, such as instruments, digital repositories, desktop Grids and virtualization. The Resource Infrastructure is:

- *integrated* if it is not provided by a EGI-InSPIRE partner, but it relies on EGI operational services (for example, the IGALC Resource Infrastructure in South America);
- *peer Grid* if it is accessible to EGI users, but it does not rely on EGI operational services (e.g. the Open Science Grid in the USA).

The current status of resource integration is documented in [MS407].

The Resource Infrastructure Provider is the legal organisation that is responsible of establishing, managing and of operating directly or indirectly the operational services to an agreed level of quality needed by the Resource Centres and the user community. It holds the responsibility of integrating them in EGI to enable uniform resource access and sharing for the benefit of their consuming end-users. The Resource Infrastructure Provider liaises locally with the Resource Centres.

For each Resource Infrastructure Provider the operational services needed for the seamless integration of different infrastructures are delivered by an Operations Centre. Locally, these are responsible for supporting their Resource Centres, monitoring their performance, collecting requirements and for representing them in the various EGI operations boards. Globally, the

Operations Centre is in charge of contributing to the development of the EGI operations roadmap and the evolution of EGI operations.

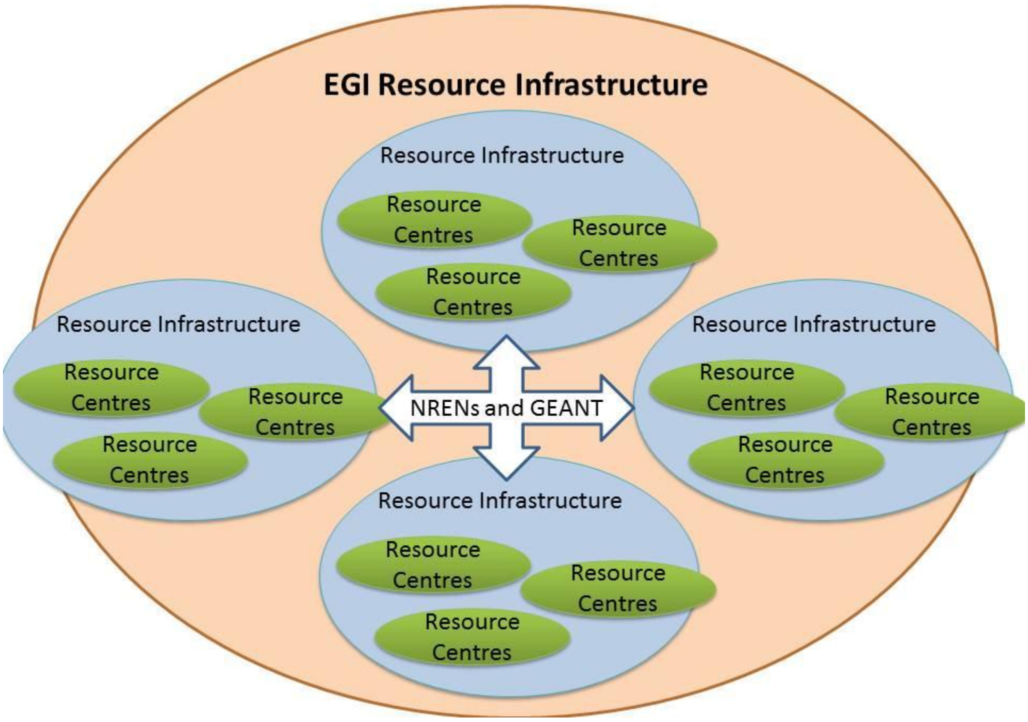

Figure 1. Architecture of the EGI-InSPIRE Resource Infrastructure.

Section 2 provides information about the number of Resource Centres, their geographical distribution across the Resource Infrastructures provided by EGI-InSPIRE members and the integrated ones provided by partners outside of the project. The operational structure of the infrastructure has been radically evolving from May 2010. The installed capacity (compute and storage) is assessed in section 3, where we illustrate the increase trend in terms of logical CPUs and HEP-SPEC 06 since the beginning of the project, and we compare it with the last year of EGEE-III. The infrastructure performance expressed in terms of availability, reliability and suspension is detailed in section 4, where the overall EGI monthly performance trend is analysed and correlated with the performance of the individual Resource Infrastructures. The resource utilization across different scientific disciplines is illustrated in section 5, which is followed by an in-depth analysis of the workload distribution among different Resource Infrastructures and active VOs (section 6). Finally, section 7 and 8 detail the current level of deployment of different middleware stacks, the distribution of core middleware services, and the status of the Staged Rollout Infrastructure needed for the validation of new software updates. Section 9 concludes the document.

2 RESOURCE INFRASTRUCTURE

2.1 Resource Centres

Table 1. EGI-InSPIRE Resource Centres and corresponding target.

Resource Centres	March 2011	Y1 Target
EGI-InSPIRE partners	332	300

The Resource Centre (also known as "site") is the smallest resource administration domain in EGI. It can be either localized or geographically distributed. It provides local resources (e.g. compute and storage) and the Grid functional capabilities necessary to make these resources accessible to authorized users such as Authentication and Authorization, Information Discovery, Data Access, etc. [UMD].

The integration of heterogeneous resources is facilitated by the adoption of common (eventually standard) interfaces to ensure that the resource can be *managed*, *monitored*, *accounted*, *supported* and properly *operated* [MS407].

At the end of EGEE-III (April 2010) the EGI Infrastructure included 323 Resource Centres, and the trend since then is illustrated in Figure 2 (a).

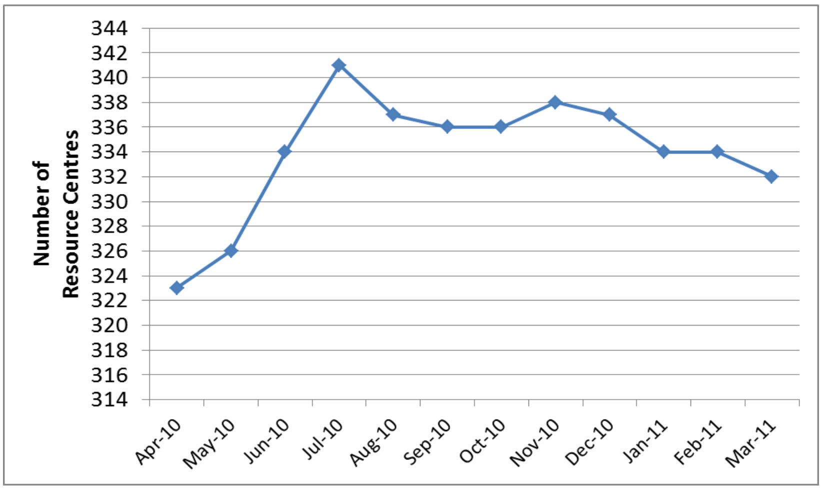

(a)

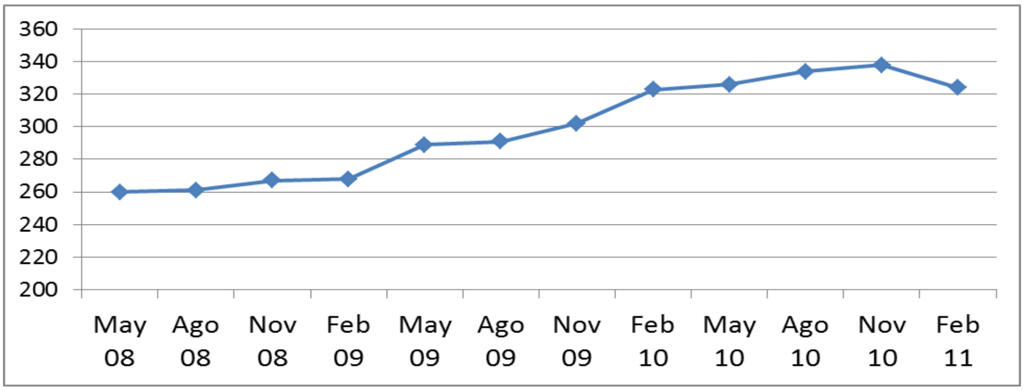

(b)

Figure 2. Certified Resource Centres that are part of the production infrastructure from May – PM1 of EGI-InSPIRE (a) and from May 2008 – PM1 of EGEE-III) (b).

The long-term trend illustrated in Figure 2 (b) shows that the Resource Infrastructure has been constantly linearly increasing without being affected by the transition from the EGEE-III project to the EGI-InSPIRE project. Short-term fluctuations in the Resource Infrastructure topology are common and are normally related to production sites that are temporarily suspended.

2.2 Countries

At the end of March 2011 the EGI Infrastructure included 332 Resource Centres distributed among 58 countries and one European Intergovernmental Research Institute - EIRO (CERN).

- 45 countries and at CERN resources are directly contributed by EGI-InSPIRE partners.
- 13 countries contribute resources through Resource Infrastructure Providers that are non-EGI-InSPIRE partners but are fully integrated with the EGI Services Infrastructure. These are: China, Pakistan and New Zealand (Asia Pacific Federation); Austria and Estonia (NDGF Federation); Belgium (Nederland Federation); Argentina, Brazil and Venezuela (GISELA Consortium); Canada and China (Canada Federation); Brazil, Chile, Colombia and Mexico (Latin America Federation).

At the end of March 2011, four EGI-InSPIRE partners are not currently contributing resources in four countries: Albania, Moldova, Indonesia and Singapore.

Additional resources are contributed through non-integrated EGI Resource Infrastructure Providers: Brazil and United States of America (OSG) and South Africa (SAGrid). These infrastructures interoperate with EGI but are not direct customers of EGI operational services.

During the last 12 month the number of countries increased from 48 at the end of EGEE-III to 58, this amounting to an increase of 20.8 % driven by the integration of new infrastructures in the Baltic and South East Europe regions. Albania and Moldova are planning to contribute resources during 2011.

The distribution of Resource Centres per country and Resource Infrastructure providers is summarized in Table 2[3].

Table 2. Distribution across European and non-European countries of EGI Resource Infrastructure Providers (EGI-InSPIRE, non-EGI-InSPIRE and peer Grids). Uncertified Resource Centres are not considered. Data source: Accounting Portal and GStat, March 2011.

Resource Infrastructure Providers		EGI Operations Centre	Resource Centres
EGI-InSPIRE Partners	Albania	NGI_AL	-
	Armenia	NGI_ARMGRID	5
	Australia	Asia Pacific Federation	2
	Belarus	NGI_BY	4
	Bosnia and Herzegovina	NGI_BA	1
	Bulgaria	NGI_BG	9
	CERN	CERN	1
	Cyprus	NGI_CYGRID	2
	Croatia	NGI_HR	3
	Czech Republic	NGI_CZ	2
	Denmark	NGI_NDGF Federation	1
	Finland	NGI_NDGF	2
	France	NGI_FRANCE	17
	FYR of Macedonia	NGI_MARGI	1
	Germany	NGI_DE	17
	Georgia	NGI_GE	1
	Greece	NGI_GRNET	13
	Hungary	NGI_HU	5
	India	Asia Pacific Federation	1
	Indonesia	Asia Pacific Federation	-
	Ireland	UKI Federation	6
	Israel	NGI_IL	3
	Italy	Italy	57

[3] The distribution of sites per country is based on the GStat tool, which extracts data from the Information Discovery System. Inaccuracies in the table are possible, if the “country” information is not published into the Information Discovery System by the Resource Centres.

	Latvia	NGI_NDGF Federation	3
	Lithuania	NGI_NDGF Federation	3
	Japan	Asia Pacific Federation	3
	Malaysia	Asia Pacific Federation	4
	Moldova	NGI_MD	-
	Montenegro	NGI_ME	1
	Netherlands	NL Federation	15
	Norway	NGI_NDGF Federation	1
	Philippines	Asia Pacific Federation	1
	Poland	NGI_PL	7
	Portugal	Ibergrid Federation	7
	Romania	NGI_RO	8
	Russia	Russian Federation	11
	Serbia	NGI_AEGIS	6
	Singapore	Asia Pacific Federation	-
	Slovakia	NGI_SK	4
	Slovenia	NGI_SI	2
	South Korea	Asia Pacific Federation	2
	Spain	Ibergrid Federation	20
	Sweden	NGI_NDGF Federation	2
	Switzerland	NGI_CH	4
	Taiwan	Asia Pacific Federation	6
	Thailand	Asia Pacific Federation	2
	Turkey	NGI_TR	6
	Ukraine	Russian Federation	2
	United Kingdom	UKI Federation	19
Non-EGI-InSPIRE partners - Integrated	Argentina	IGALC Federation	1
	Austria	NGI_NDGF Federation	2
	Belgium	NL Federation	4
	Brazil	LA Federation (2) IGALC Federation (2)	4

	Canada	Canada Federation	7
	Chile	LA Federation	1
	China	Canada Federation (2) Asia Pacific (1)	3
	Colombia	LA Federation	1
	Estonia	NGI_NDGF	2
	Mexico	LA Federation	1
	New Zealand	Asia Pacific Federation	1
	Pakistan	Asia Pacific Federation	1
	Venezuela	IGALC Federation	1
Peer non-European Grids	OSG (Brazil)	GridUNESP_CENTRAL SPRACE (OSG)	2
	OSG (USA)	OSG	OSG (USA)
	SAGrid (South Africa)	SAGrid	4

2.3 Resource Infrastructure Providers

EIROs and the National Grid Initiatives (NGIs) are Resource Infrastructure Providers. In Europe Resource Centres are required to be affiliated to the respective NGIs, which (a) have a mandate to represent their national Grid community in all matters falling within the scope of EGI.eu, and (b) are the only organization having the mandate described in (a) for its country and thus provide a single contact point at the national level.

In addition, EGI integrates with a number of external Resource Infrastructure Providers, where *external* qualifies those providers that either have no representation in the EGI Council or no partnership in the EGI-InSPIRE project. In these cases, the framework of collaboration is defined through the *Resource Infrastructure Provider MoU* [MoU] which aims at allowing an integrated access to Resource Infrastructures, sharing operational services, committing to a common set of policies and procedures, and cooperating for the evolution of a common operations architecture. The Resource Infrastructure Provider MoU is currently under negotiation with the GISELA Consortium (representing CEDIA, CEFET-RJ, CEFET-RJ, UFCG, UFRJ, ULA-MERIDA and USB) and the South African National Grid (SAGrid4).

At the end of March 2011 EGI comprised 40 NGIs and one EIRO (CERN) of which two (Albania and Moldova) are expected to join during 2011.

[4] http://www.sagrid.ac.za/

2.4 Operations Centres

At the end of the EGEE-III project the Resource Infrastructure was operated by 14 Regional Operational Centres (ROCs): Asia Pacific, Canada, Central Europe, CERN, France, Germany/Switzerland, IGALC, Italy, Latin America, Northern Europe, Russia, South Eastern Europe, South Western Europe, and United Kingdom/Ireland. This scenario has evolved considerably during the first project year of EGI-InSPIRE. The largest ROCs (Central Europe and South East Europe) stopped their operations during PQ2and PQ3 respectively. The EGEE ROCs have consequently developed into a much larger group of smaller Operations Centres [D4.1], which typically serve a single country.

There are 32 EGI Operations Centres in total.

Currently the 40 NGIs are operated by 27 Operations Centres. The overall current state is illustrated in the following section.

- **European national Operations Centres:** 22 active centres operated by NGIs at a national level in the following countries: Armenia, Belarus, Bosnia and Herzegovina, Bulgaria, Cyprus, Croatia, Czech Republic, France, FYR of Macedonia, Germany, Georgia, Greece, Hungary, Israel, Italy, Montenegro, Poland, Romania, Serbia, Slovakia, Slovenia and Turkey.
- **European federated Operations Centres:** five multi-country centres providing services to 16 NGIs in total.
    1. *IberGrid*: Portugal and Spain;
    2. *Netherlands Federation*: Belgium and Netherlands;
    3. *Russian Federation*: Russia and Ukraine;
    4. *NDGF Federation*: Austria, Denmark, Estonia, Finland, Latvia, Lithuania, Norway and Sweden;
    5. *United Kingdom/Ireland Federation*: Ireland and United Kingdom (decommissioning is planned for 2011).
- **EIRO Operations Centre**: CERN
- **Non-European Operations Centres**: only federated centres are currently active.
    1. *Asia Pacific Federation*: Australia, China, India, Japan, Malaysia, Philippines, South Korea, Taiwan and Thailand;
    2. *Canada Federation*: Canada and China;
    3. *GISELA Consortium* (IGALC Federation): Argentina, Brazil, and Venezuela;
    4. *Latin America*: Brazil, Chile, Colombia and Mexico.

The commissioning/decommissioning chronology of the EGI federated Operations Centres from May 2010 until March 2011 is illustrated in Figure 3.

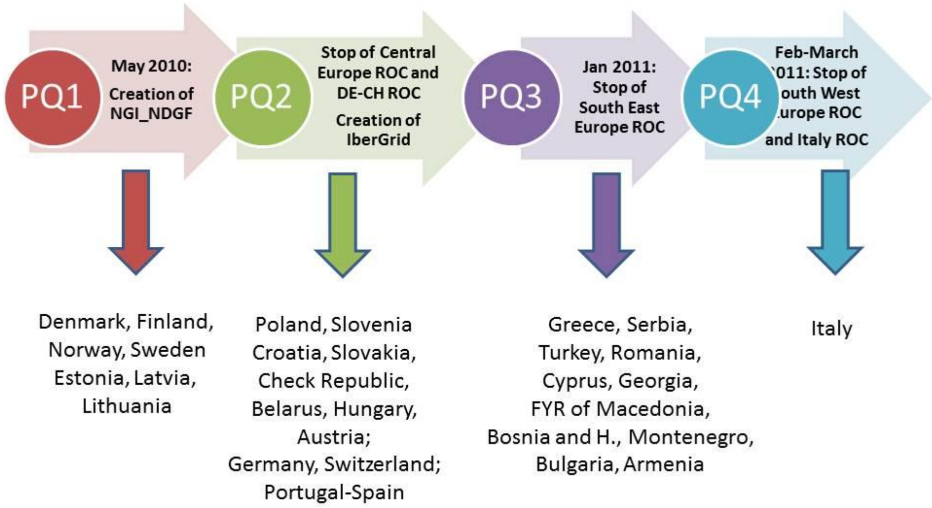

Figure 3. Commissioning/decommissioning chronology of the EGI federated Operations Centres from May 2010 until March 2011.

The distribution of Resource Centres among the active Operations Centres is illustrated in Figure 4. As the diagram shows, only 4 Operations Centres operated more than 20 centres; these are: Asia Pacific, IberGrid, Italy and and United Kingdom/Ireland.

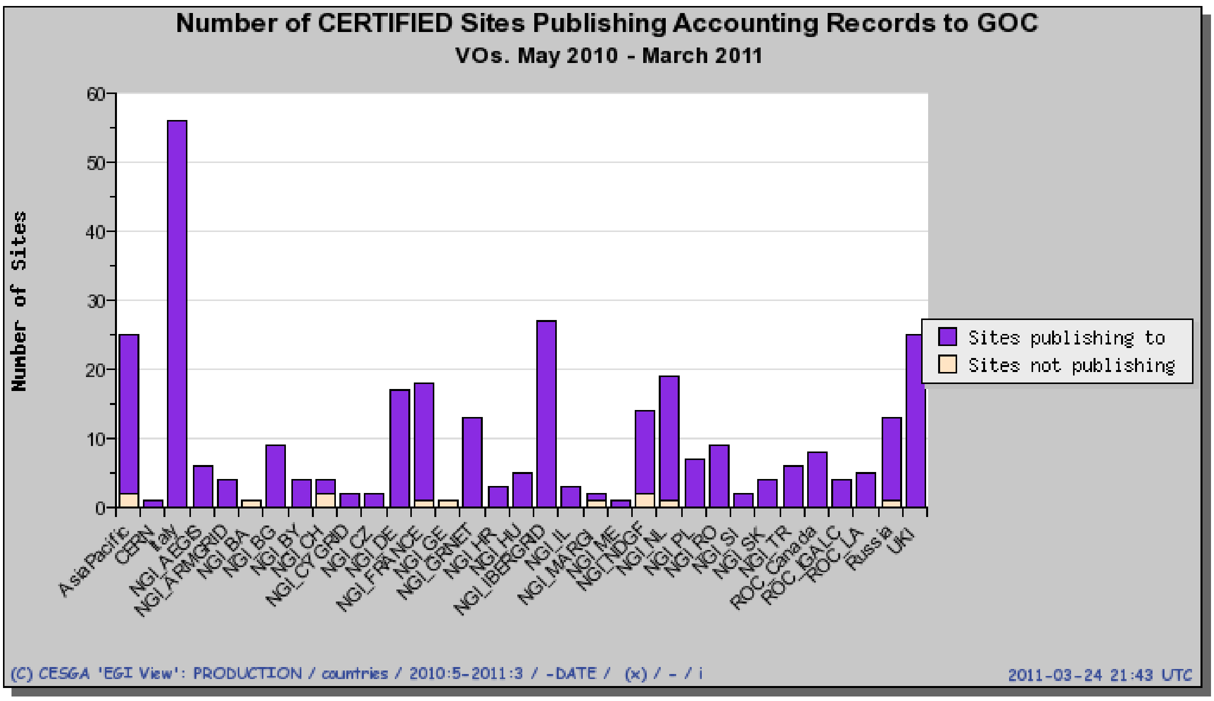

Figure 4. Distribution of Resource Centres amongst the EGI Operations Centres (May 2010 – March 2011). Source: EGI Accounting Portal.

3 INSTALLED CAPACITY

3.1 Compute Resources

The target amount of logical CPUs (cores) expected for the first year of the project amounts to 200,000 units, this only includes the resources contributed by EGI-InSPIRE partners and thus excluding logical CPUs contributed by:

- integrated Resource Infrastructures Providers (those that are currently using EGI operational services but are not a EGI-InSPIRE);
- peer non-European Grids such as OSG and South Africa Grid.

Table 3 indicates the estimated compute capacity installed capacity in Project Quarter (PQ) 1, 2 and 3, the amount of contributed logical CPUs met the project target value of 200,000 units in production, and has been steadily increasing since March 2010. As detailed in Table 3, the increase from April 2010 – at the end of the EGEE-III – amounts to 7.9 % for logical CPUs and 39.8 for HEP-SPEC 06.

HEP-SPEC06 is the EGI reference performance benchmark of compute resources [HS06]. It was defined by the HEPiX Benchmarking Working Group [HWG] is and it based on SPEC. One HEPSPE06 corresponds approximately to 250 SI00 (this was tested with) HEP applications, and in what follows the normalized CPU time consumed is chosen as reference metric (HEP-SPEC06 Hours).

The project quarterly reports were used as authoritative source of information for all partners that contributed feedback. In case of missing information data was collected from GStat[5]. GStat is a visualization tool that extracts information from top-level Information Discovery Systems. As stated in the Grid Site Operations Policy (point 3) [SOP], it is a responsibility of the Resource Centre administrators to ensure that publishes information is accurate. However publishing of information is error prone, as installed capacity is in some cases is manually configured and published data can be consequently inaccurate.

Table 3. EGI-InSPIRE logical CPUs

Logical CPUs	PQ1	PQ2	PQ3
EGI-InSPIRE partners (Y1 target: 200,000)	184,844	197,777	207,203
EGI-InSPIRE partners and integrated infrastructures	277,193	296,588	308,583

[5] http://gstat.egi.eu/gstat/geo/openlayers

Table 4. Installed compute capacity (logical CPUs and Million HEP-SPEC 06) in April 2010 and March 2011

	April 2010 (EGEE-III Infrastructure)[6]	March 2011 (EGI-InSPIRE Infrastructure)	Increase (%)
Logical CPUs	192,000	207,203	7.9
Million SI2k	335 (1.34 Million HEP-SPEC 06)	495 (1.98 Million HEP-SPEC 06)	47.7

The distribution of compute capacity (logical CPUs and HEP-SPEC 06) per Resource Infrastructure Provider (NGI and EIRO) that is partner of EGI-InSPIRE is illustrated in Figure 5 (a).

The overall installed compute capacity including the EGI integrated Resource Infrastructures is illustrated in Figure 5 (b).

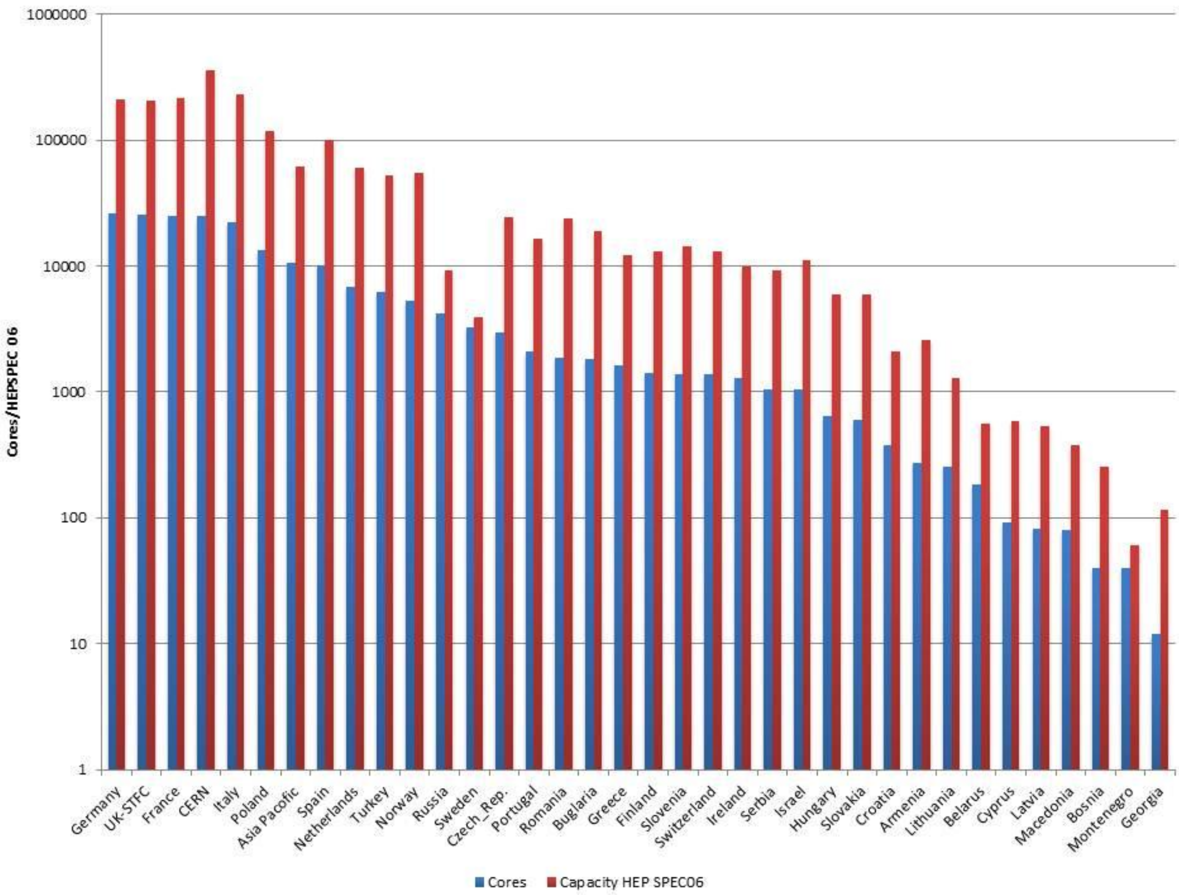

[6] Data source: EGEE Deliverable DSA1.2.2 (https://edms.cern.ch/file/1060571/3/EGEE-III-DSA1.2.2-1060571-v2.pdf)

(a)

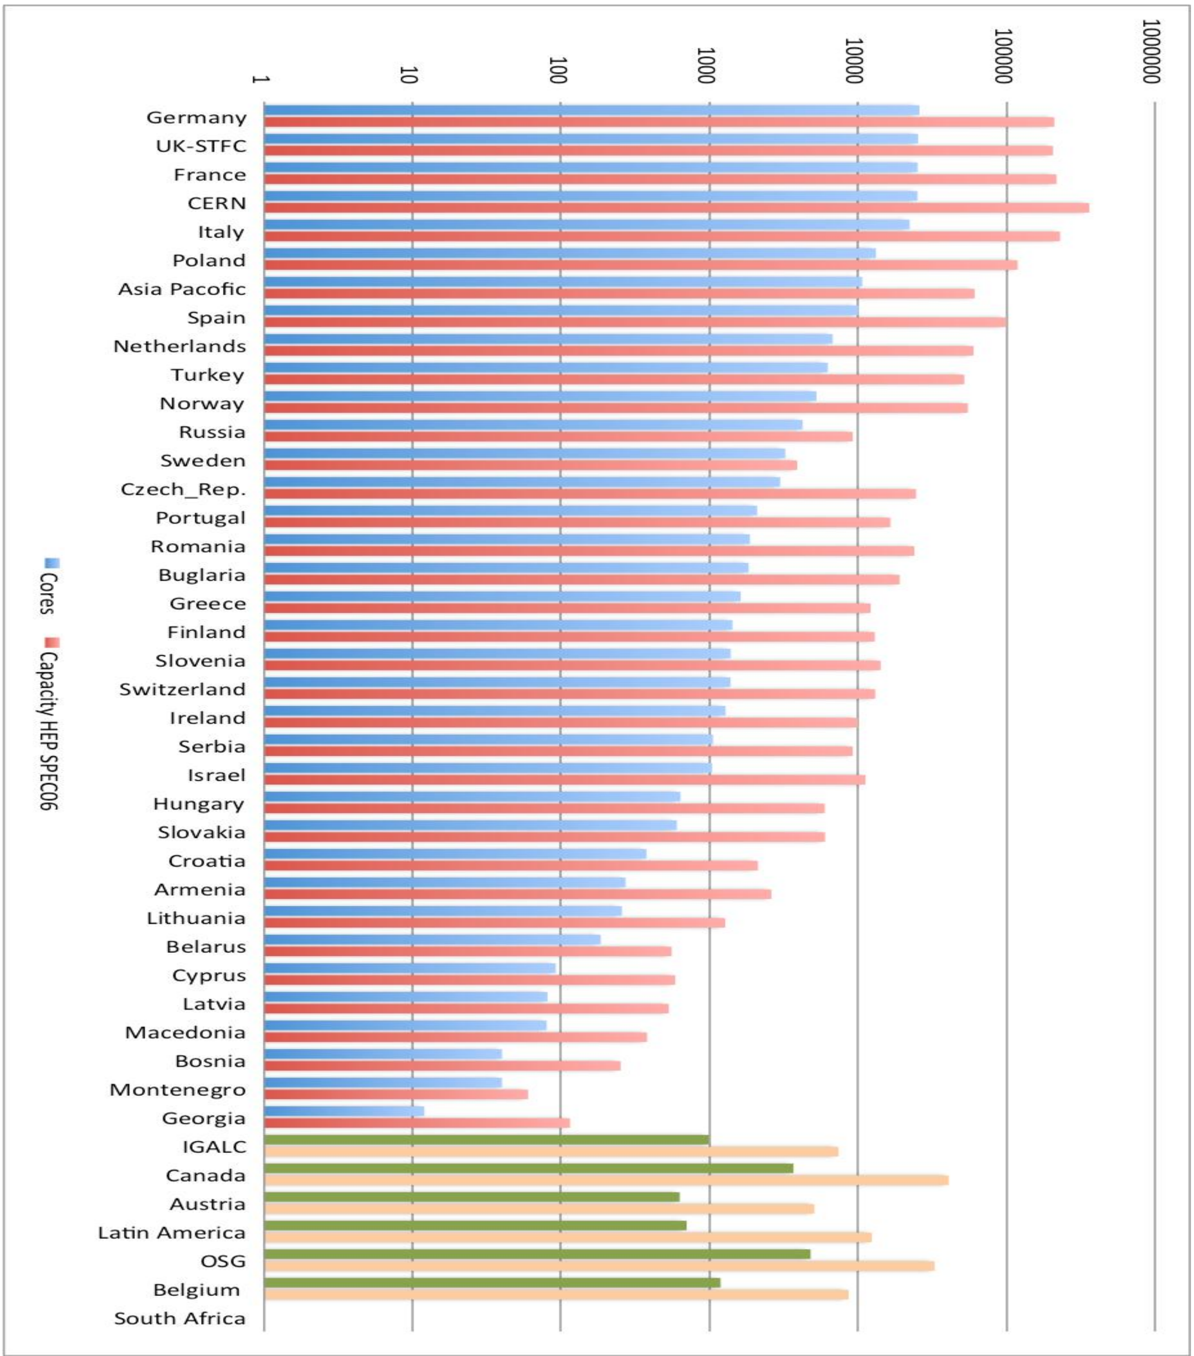

(b)

Figure 5. Distribution of compute installed capacity (cores and HEP-SPEC 06) (a) per Resource Infrastructure Provider (NGIs and EIROs) and (b) including integrated Resource Infrastructures – March 2011 (source: project quarterly reports and GStat).

3.2 Storage Resources

Information about storage capacity provided by the EGI-InSPIRE partners is periodically collected through the project quarterly reports. For EGI-InSPIRE in case of missing input such information is not available, the GStat tool is used as source. 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.

The total amount of reported installed disk capacity amounts to 101 PB (40 PB at the end of EGEE-III). The distribution of disk storage resources among the EGI-InSPIRE partners is illustrated in Figure 6.

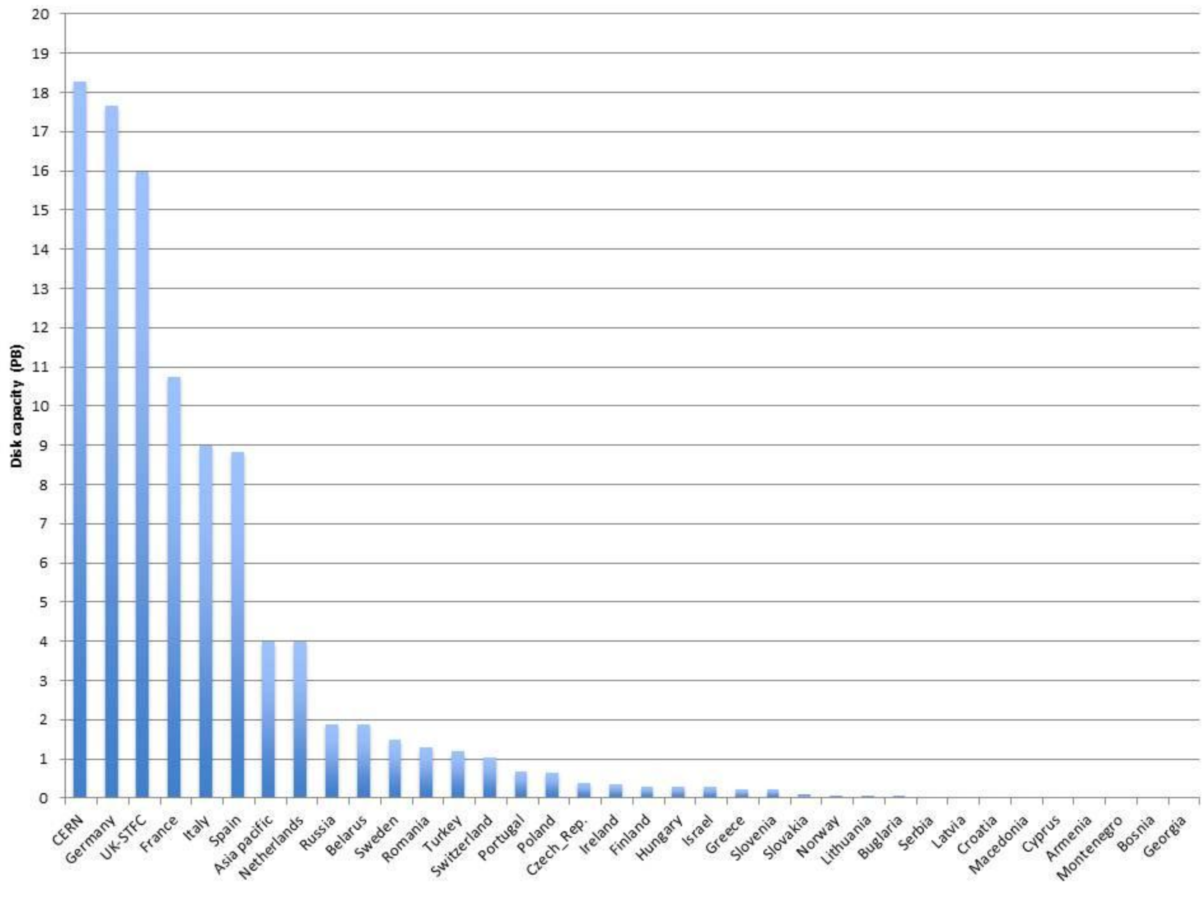

Figure 6. Installed disk capacity in PB across the EGI-InSPIRE partners – March 2011 (source: project quarterly reports and GStat).

Tape capacity is mainly provided by Resource Centres contributing resources to the LHC experiments. The total installed tape (also known as *nearline*) capacity amounts to 80 PB at the end of March 2011 as reported by GStat[7].

[7] http://gstat-wlcg.cern.ch/apps/capacities/site_storage/

3.3 Compute Resources for Parallel Jobs

Table 5. Integration metrics (HPC and MPI)

Metric	Y1 Target	PQ3
Number of HPC clusters (M.SA1.Integration.1)	1	54
Number of sites with MPI (M.SA1.Integration.2)	50	90

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. Myrinet[8], InfiniBand[9]). The clusters that qualify as high-performance, have been reported by the Resource Infrastructure Providers to amount in total to 54 units at the end of PQ3 (see Table 5).

At the end of PQ3 Message Passing Interface [MPI] jobs were supported by 90 Resource Centre. At the end of PQ2 the overall amount of Resource Centres supporting MPI amounted to 21.66 % and this increased to 26.47 % at the end of PQ3 (Figure 7).

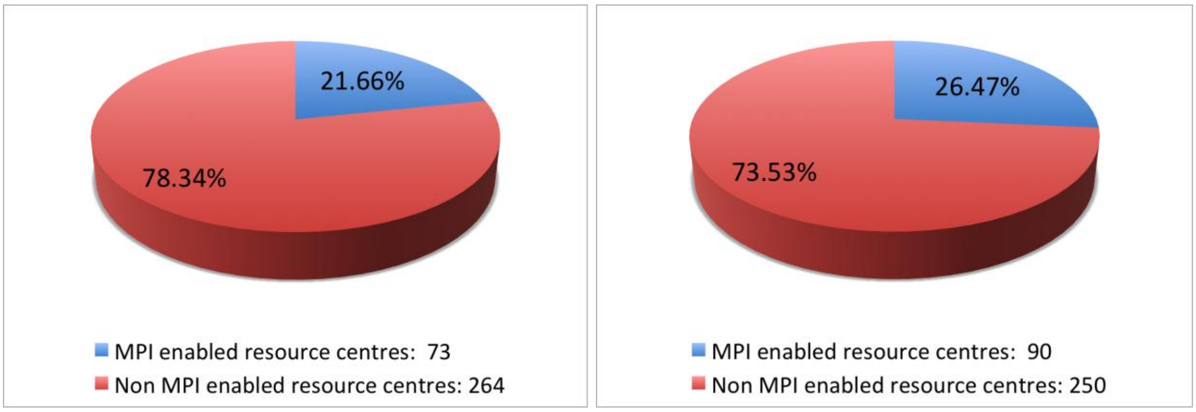

(a) (b)

Figure 7. Number of Resource Centres supporting MPI jobs in PQ2 (a) and PQ3 (b) (source: project quarterly reports and GridMap)

Testing of MPI job submission is currently part of the standard set of Nagios probes that are deployed by the Resource Infrastructure Providers. Results of MPI tests are daily monitored by operations personnel on duty and failures produce alarms in the Operations Dashboard[10], for a standard and more effective control and support of MPI in the infrastructure.

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

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

[10] https://wiki.egi.eu/wiki/SAM_Tests#Operations_tests

4 PERFORMANCE

4.1 Availability and Reliability

Table 6. Availability and Reliability metrics.

EGI Average Monthly Reliability	May 2010-March 2011	Y1 Target
Reliability	91.9 %	90 %
Availability	90.7 %	-

Availability and Reliability are the two metrics that EGI-InSPIRE utilizes to measure the quality of operational services delivered by Resource Centres, Resource Infrastructures and EGI. These are computed by collecting the results of the periodic tests performed for all certified centres through the OPS VO, which is dedicated to monitoring activities. OPS tests provide a good indication of the overall performance of the operations of a Resource Centre. For measurement of the Availability and Reliability perceived by users, VO-specific tests need to be performed, which are customized according to the VO workflow and computing model. The remainder of the chapter focuses on OPS statistics.

Availability and Reliability statistics are computed by collecting and summarizing monitoring results on hourly, daily, weekly and monthly basis (currently through the GridView tool, and in the future through the Availability Computation Engine – a new software component currently under validation and developed in the framework of the WLCG collaboration).

The overall Availability and Reliability of a Resource Infrastructure depend on the individual performance of the related Resource Centres. As different centres provide different amounts of installed resource capacity, the overall Availability and Reliability of the Resource Infrastructure are expressed as the weighted average of the individual Resource Centre figures, where the weight is proportional to the amount of compute capacity contributed by the site expressed in HEP-SPEC 06[11]. The EGI Availability and Reliability were computed as the arithmetic mean of the weighted Availability and Reliability of all certified sites. This is illustrated in Figure 8, were the EGI overall monthly Availability and Reliability are compared. From May 2011 Figure 8 (a) the monthly Availability and Reliability have been fluctuating in the range [84.4, 95.85] %, without a clear increase trend. The average monthly EGI Availability and Reliability respectively amount to 90.7 % and 91.9 %. Both values are in-line with the project target Reliability of 90 % and with the performance at the end of EGEE-III.

When looking at the Availability and Reliability from Jan 2009 Figure 8 (b) we can see that monthly performance has been fluctuating during the second year of EGEE-III too. Both the last year of EGEE-III and the first of EGI-InSPIRE are characterized by a radical change in the structure of the active Operational Centres and oscillating performance is partially due to this. During Year 2 three new Operations Centres were created: Canada, Latin America, and IGALC. Similarly, in the first year of EGI-InSPIRE the largest federated Operations Centres dissolved to generate a set of smaller ones which started their own independent operational activities (section 2.4). In particular, the lowest

[11] Running HEP-SPEC: https://twiki.cern.ch/twiki/bin/view/FIOgroup/TsiBenchHEPSPEC

EGI-InSPIRE performance, which was scored in June 2010 and January 2011, correspond with the end of operations of Central Europe ROC and South East Europe ROC. In these two months many new Operations Centres became operational, namely Slovenia, Croatia, Slovakia, Check Republic and Serbia in June 2010, and Bosnia Herzegovina, FYR of Macedonia, Montenegro, Bulgaria and Armenia in January 2011. The individual NGI monthly performance is plotted in Figure 9.

Other factors related to the current algorithm for availability computation have in some cases caused low performance. One of these is the process of certification of new sites. As currently the topology information system does not provide historical information, if a Resource Centre is "certified" after the end of the month but before the availability statistics are calculated, then the newly certified site is considered to be "certified" also during the reference month, and its availability statistics are included even if during the reference month the Resource Centre was uncertified. This can contributed to lower the performance of an NGI, especially if the number of sites under certification is considerable, or if the fraction of resources contributed by the newly certified is large.

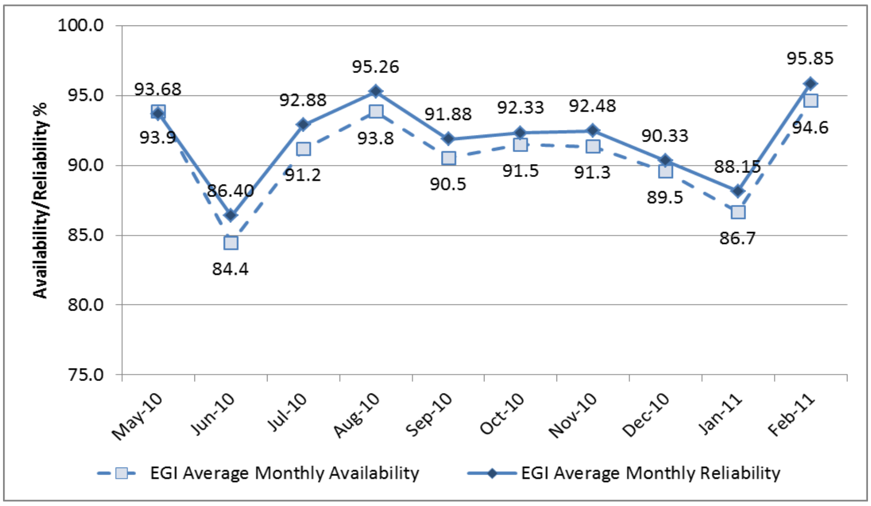

(a)

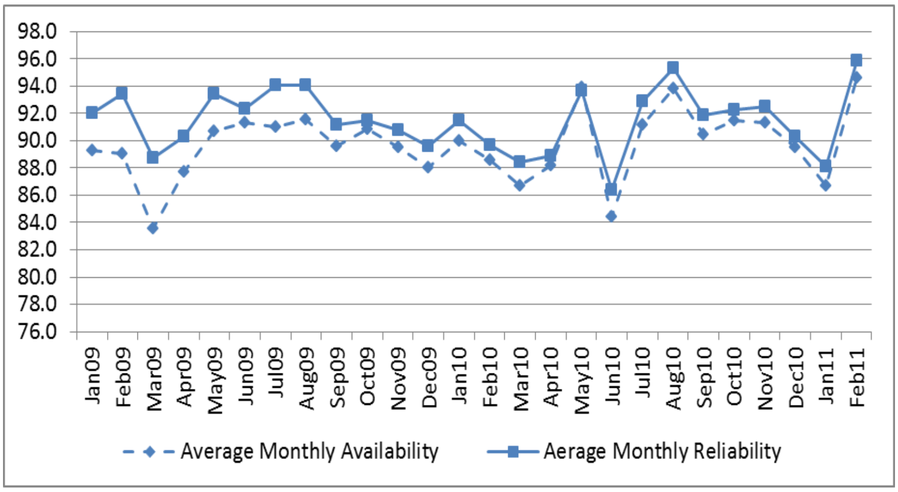

(b)

Figure 8: Total EGI-InSPIRE Availability and Reliability trend from May 2010 to March 2011 – EGI-InSPIRE (a) and from January 2009 to March 2011 – EGEE and EGI-InSPIRE (b)

Infrastructure performance has been monitored since the beginning of the project. In May 2010 the process and responsibilities for the distribution of the reports was revised. In parallel, the procedure to collect justifications for underperforming sites was revised and streamlined: currently all justifications from Resource Centres are handled in GGUS tickets and are collected centrally by the COD team[12], who are also responsible of following-up suspension cases.

Several actions will be undertaken during 2011 to mitigate the problem of underperforming centres:

- The GOCDB repository will provide time information about changes of status of the resource centres, in this way historical information can be retrieved by the Aggregated Topology Provider[13] for an increased accuracy of the monthly statistics.
- A notification system is needed to warn the administrators of the underperforming centres before the end of the month, in order to allow them to adopt counter measures and mitigate the problem. A Nagios-based notification system is under discussion. In addition, information about the running monthly availability and reliability should be available from the Operations Portal. In case of low performance alarms can be generated, which can be used to open a tickets to the affected sites. The same tickets can be used to collect information on performance issues, thus streamlining the performance follow-up procedure of COD.

Individual NGI statistics are plotted in Figure 9. The complete repository of the infrastructure performance statistics from January 2009 is accessible from the EGI wiki [AVA].

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

[13] https://wiki.egi.eu/wiki/SAM#ATP

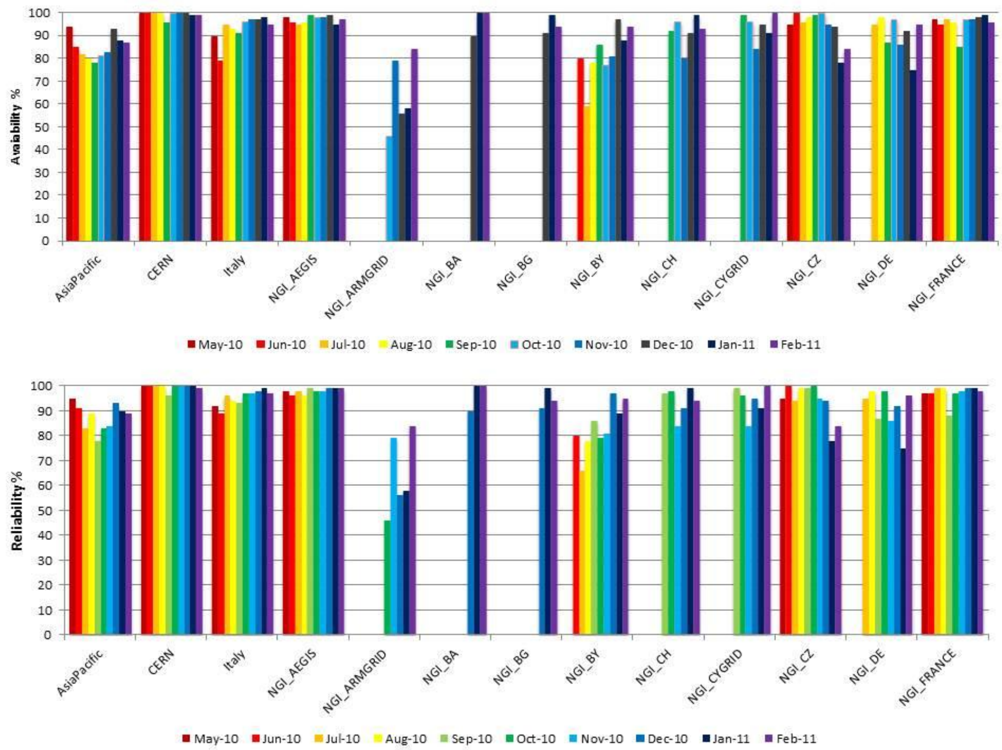
Availability %
AsiaPacific
CERN
Italy
NGI_AEGIS
NGI_ARMGRID
NGI_BA
NGI_BG
NGI_BY
NGI_CH
NGI_CYGRID
NGI_CZ
NGI_DE
NGI_FRANCE
May-10 Jun-10 Jul-10 Aug-10 Sep-10 Oct-10 Nov-10 Dec-10 Jan-11 Feb-11
Reliability %
AsiaPacific
CERN
Italy
NGI_AEGIS
NGI_ARMGRID
NGI_BA
NGI_BG
NGI_BY
NGI_CH
NGI_CYGRID
NGI_CZ
NGI_DE
NGI_FRANCE
May-10 Jun-10 Jul-10 Aug-10 Sep-10 Oct-10 Nov-10 Dec-10 Jan-11 Feb-11

(a)

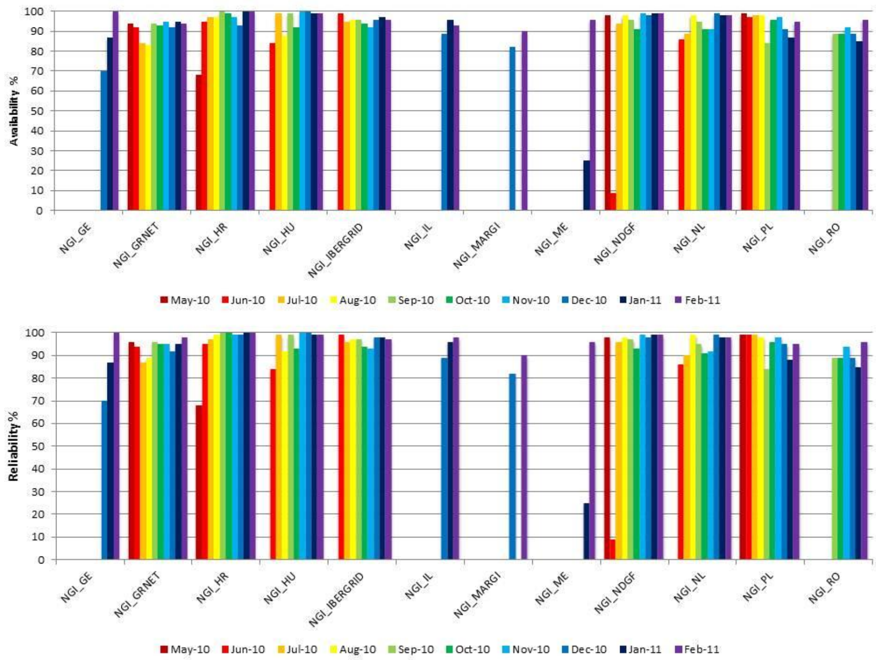
Availability %
NGI_GE
NGI_GRNET
NGI_HR
NGI_HU
NGI_IBERGRID
NGI_IL
NGI_MARGI
NGI_ME
NGI_NDGF
NGI_NL
NGI_PL
NGI_RO
May-10 Jun-10 Jul-10 Aug-10 Sep-10 Oct-10 Nov-10 Dec-10 Jan-11 Feb-11
Reliability %
NGI_GE
NGI_GRNET
NGI_HR
NGI_HU
NGI_IBERGRID
NGI_IL
NGI_MARGI
NGI_ME
NGI_NDGF
NGI_NL
NGI_PL
NGI_RO
May-10 Jun-10 Jul-10 Aug-10 Sep-10 Oct-10 Nov-10 Dec-10 Jan-11 Feb-11

(b)

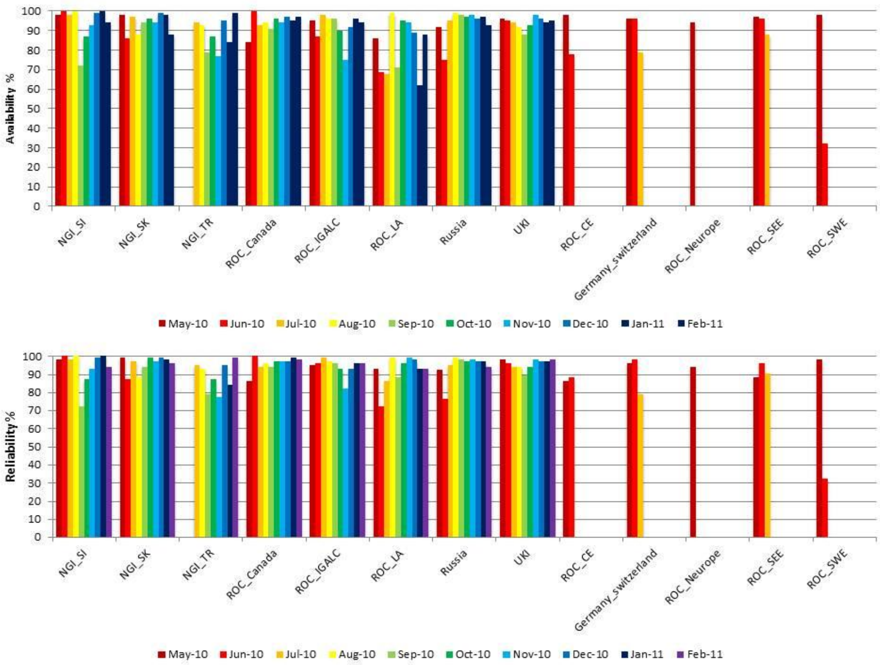

(c)

Figure 9. Monthly Availability and Reliability of EGI Resource Infrastructures grouped per Operations Centres. As some of the Operations Centres were decommissioned or started operations after May 2010, in these cases statistics can be incomplete.

4.2 Suspension

EGI Resource Centres are eligible for suspension if they meet one of the following conditions.

1. The Resource Centre has been underperforming for three consecutive months (the monthly availability is lower than 50 %) [PROC04].
2. The Resource Centre is affected by a critical vulnerability issue and no repairing action is successfully undertaken in due time [SEC03].
3. The Resource Centre is in downtime for more than one month [MAN02].

From the beginning of EGI-InSPIRE several Resource Centres were suspended, in all cases this was due to performance issues or lack of responsiveness to performance tickets: three centres from Asia Pacific Federation, two from NGI_ARMGRID, and one from the Russian Federation[14].

The list of suspended sites is available on the EGI wiki[15].

[14] Details on suspended sites are available for the EGI wiki: https://wiki.egi.eu/wiki/Underperforming_sites_and_suspensions

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

4.3 Security

From May 2010 EGI CSIRT handled in total 6 security incidents, 9 security advisories of which two were critical requiring mandatory patching in seven calendar days [SEC03]. In both cases no centres were suspended and all managed to restore the security of the centre in due time with the support of the EGI CSIRT.

This is the detailed breakdown:

- QR1: handling of three security incidents and issuing one security advisory on a vulnerability found in Intel compiler suite.
- QR2: handling of 2 security incidents, issuing of six security advisories, of which one critical, two moderate and three high.
- QR3: handling of one security incident, issuing of three security advisories on Linux vulnerabilities, of which one was critical (CVE-2010-4170) and two at high risk.

5 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.

According to the Operations Portal[16], at the end of March 2011 there were 13319 users (9.5% increase from March 2010), while the VOs was 186 (+17.4% increase from March 2010).

The number of users and VOs herein reported is a conservative estimation, as the Operations Portal retrieves user information from the instances of the VO Membership Service (VOMS) that are indicated on the VO ID cards and if these are unavailable or incorrectly registered, membership information cannot be extracted. In addition, not all VOs (especially national ones) are currently registered centrally.

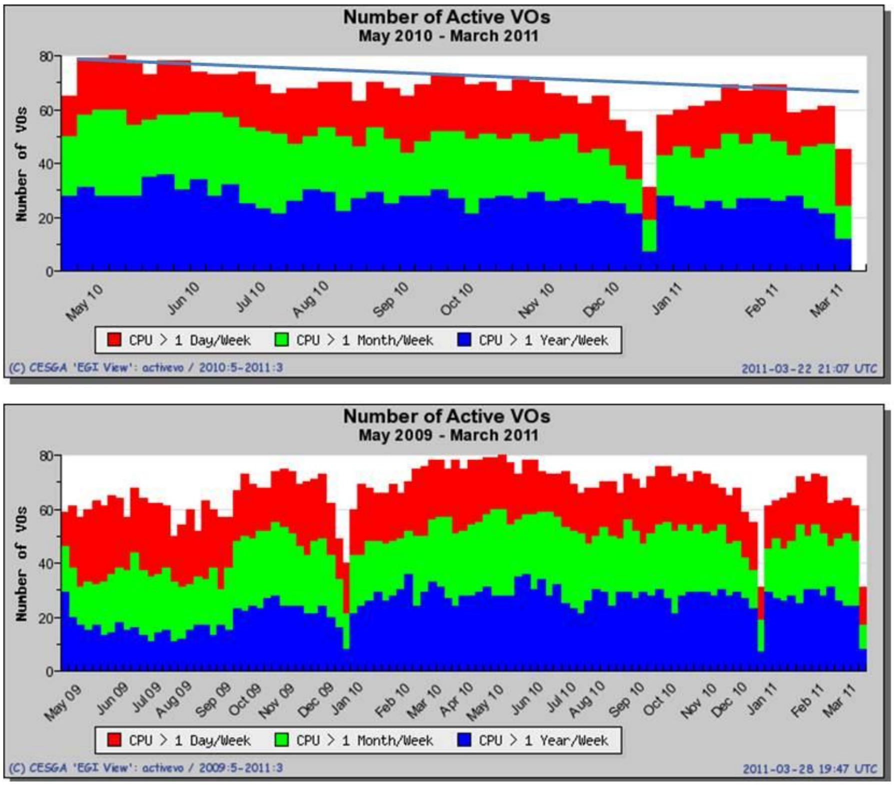

Figure 10. Number of active VOs per category (source: Accounting Portal)

For the 186 reported VOs the workload produced varies considerably (Figure 10). VOs can be differentiated into three categories according to the corresponding resource consumption: VOs with

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

low activity – red (consuming more than 1 *Day* of CPU time per week), VOs with *medium activity* – *green* (more than 1 Month of CPU time per week), and *high activity* – *blue* – consuming more than 1 *Year* of CPU time per week. As we can see in Figure 10 (b) the overall number of active VOs hasn't been increasing since May 2009. On the other hand, after a small peak from March to May 2010, the number of low activity VOs has been gradually decreasing and the total number of active VOs has stabilized around 60 units, as shown in Figure 10 (a).

5.1 VO Distribution

The distribution of registered VOs among disciplines (Figure 11) shows that the discipline with the highest number of VOs is "High-Energy Physics" (20.5%), followed by "Infrastructure" (19.2%) – it comprehends operational VOs such as OPS and DTEAM and various national catch-all VOs that include end-users, "Multidisciplinary" (14.2%) and "Astronomy Astrophysics and particle Physics" (10%). For the remaining disciplines the number of VOs is less than 10%.

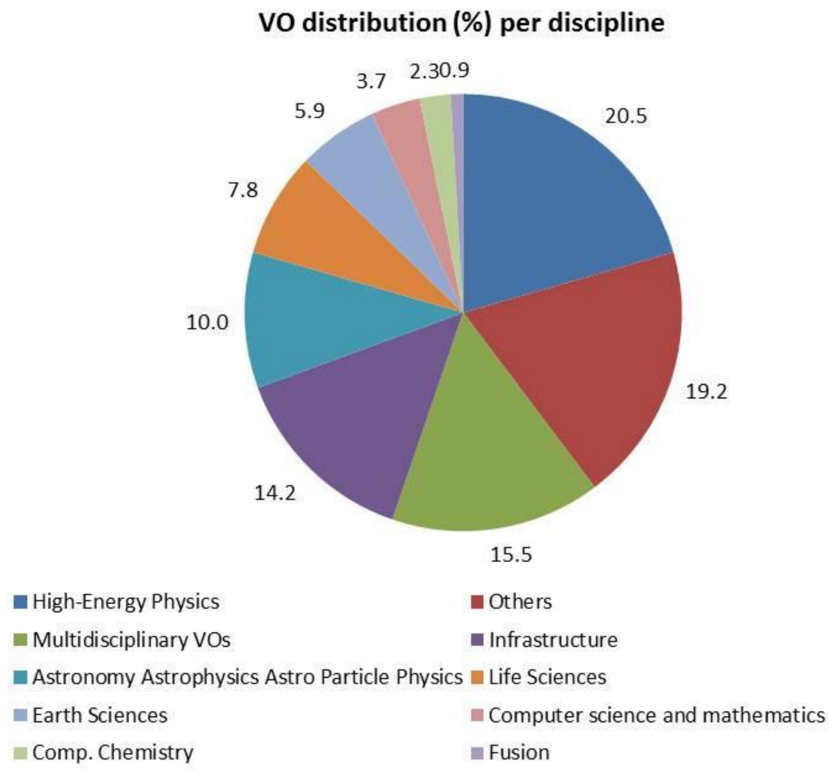

Figure 11. VO distribution per discipline (source: Operations Portal)

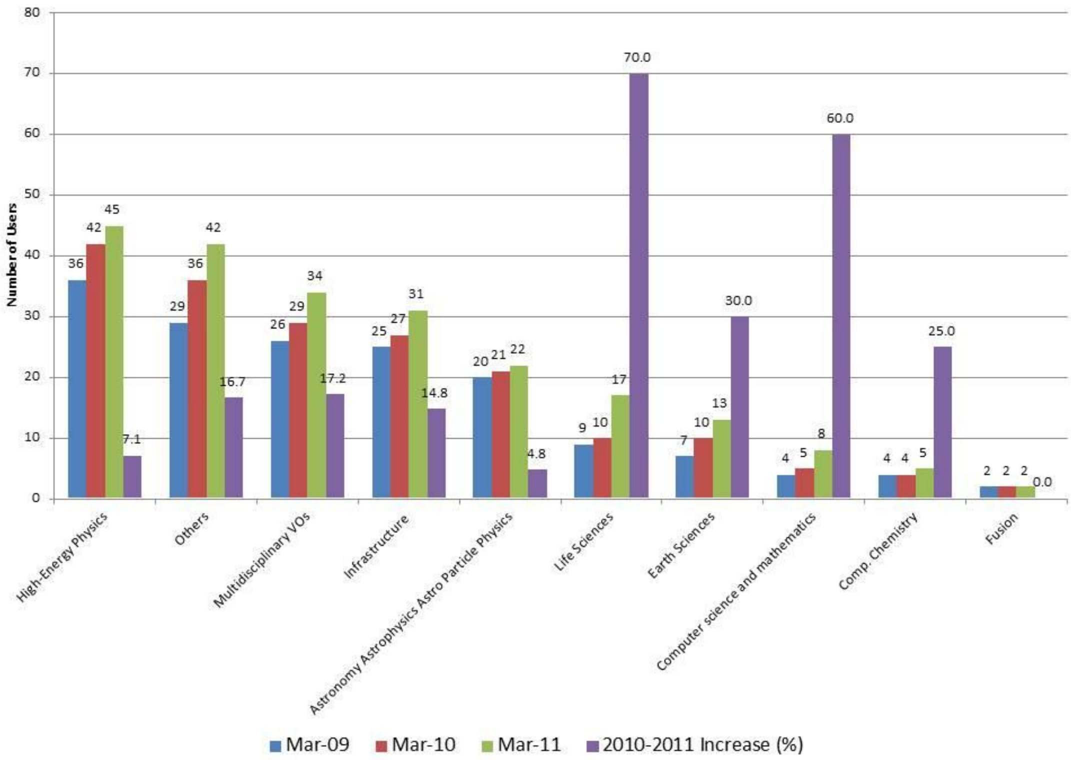

Figure 12. Number of VOs per discipline in March 2009 (blue), March 2010 (red) and March 2011 (green). For each VO the increase from 2010 to 2011 is plotted (source: Operations Portal).

As shown in Figure 12, from 2009 to date the number of VOs has been gradually increasing for all disciplines. The *relative* increase of the size of largest disciplines is small if compared to Life sciences, Earth Sciences, Computer Science and Mathematics and Computational Chemistry. The Fusion community is peculiar as the number of VOs is stable since 2009. Unfortunately, in combination with this as shown in Figure 14, the number of users has been considerably decreasing during the last 12 months (-77%). In this case, because the number of users in Fusion is not high, this decrease has no visible effect at the global level.

5.2 User Distribution

According to the data provided by the Operations Portal, users participating to the four top disciplines (High-Energy Physics, Multidisciplinary VOs, Infrastructure and Others) currently amount to 86% of the user community (see Figure 14). Among all disciplines, only in High-Energy Physics and partly "Others" the number of users has been increasing considerably during the last 12 months. For other disciplines Multidisciplinary VOs and Infrastructure the number of users decreased; for the Infrastructure discipline this can be related to the formation of new Operations Centres, to the decommissioning of existing local VOs for monitoring, or VOs that were formerly set in the framework of the BalticGrid and SEE-Grid projects. For the remaining disciplines the increase/decrease was moderate, with the only exception of Computer Science and Mathematics and Fusion, for which the relative decrease is significant (Figure 14).

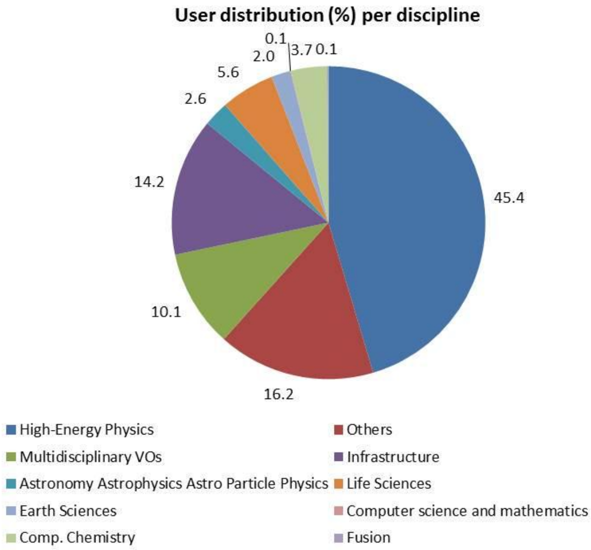

Figure13. Distribution of users among the EGI disciplines (March 2011, source: Operations Portal)

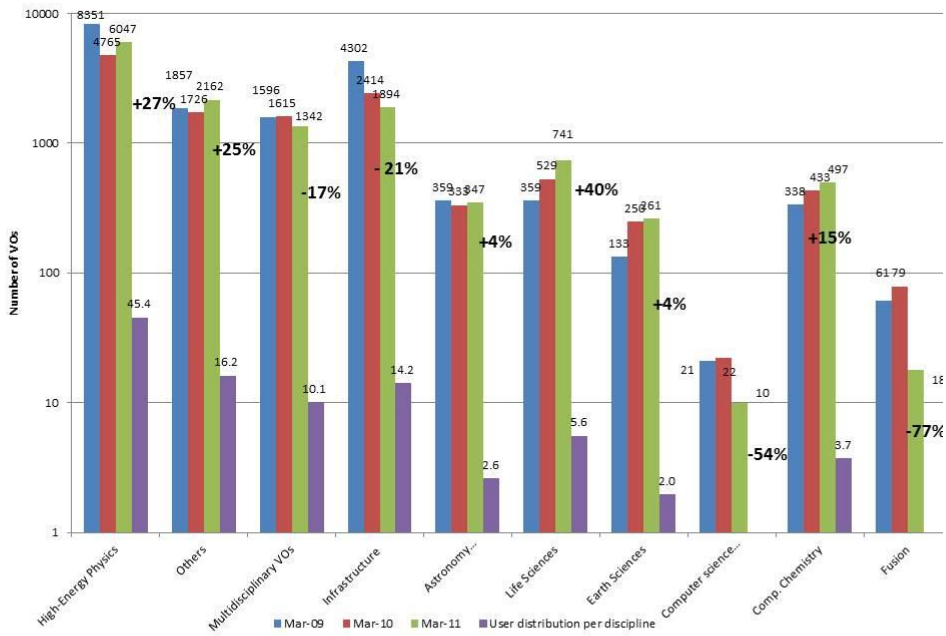

Figure 14. Increase/decrease of the number of users per discipline from March 2009 to date (logarithmic scale). The percentage in bold shows the relative increase/decrease from March 2010 to date (source: Operations Portal).

5.3 Resource Utilization per Discipline

During the first year of EGI-InSPIRE the High-Energy Physics discipline (it contributes 45.4% of the user community) has greatly expanded in resource utilization, which in total amounts to 89.9% of the overall consumed normalized CPU time measured in HEP-SPEC06, as shown in Figure 15. This trend was driven by the data taking and analysis activities of the LHC experiments. On the other hand, the cumulative utilization of the other disciplines amount to 10.1%.

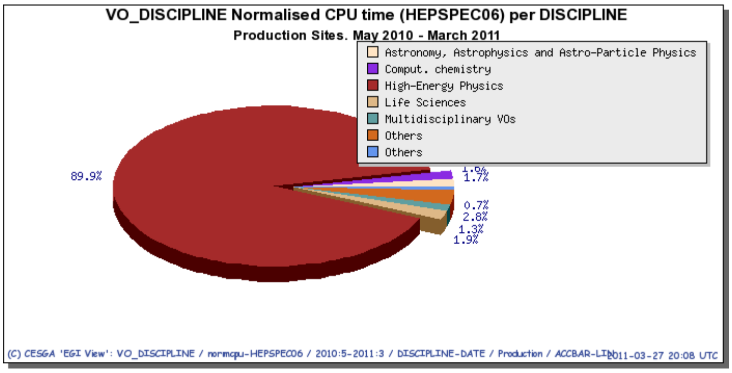

Figure 15. Resource utilization (CPU normalized time in HEP-SPEC06 hours from May 2010 to March 2011 (source: Accounting Portal).

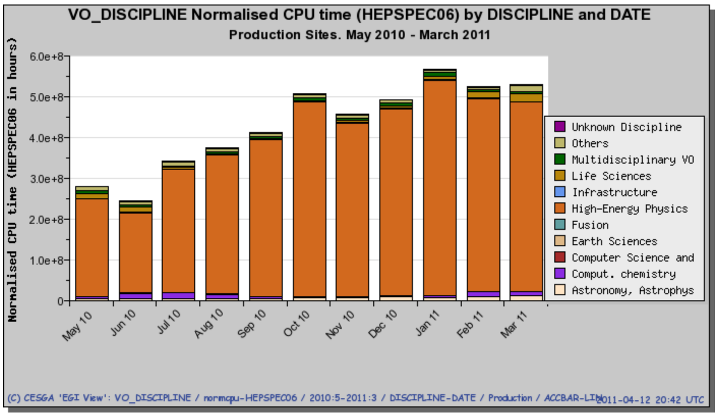

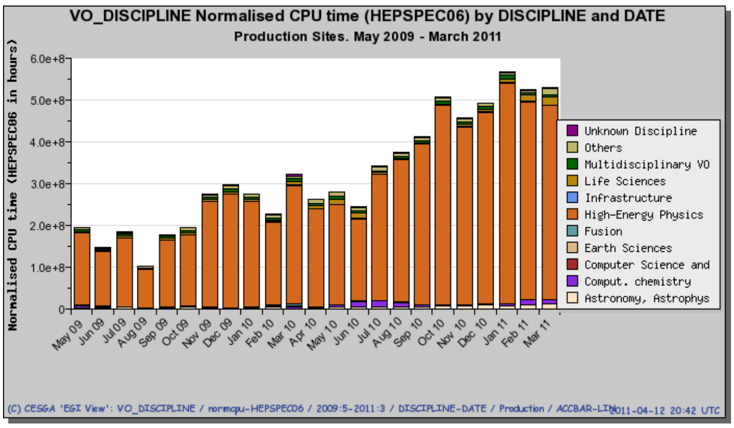

Figure 16. Normalized CPU time consumed per discipline since March 2010 (top) and since May 2009 (bottom).

As shown in Figure 16, since the beginning of the project High-Energy Physics has greatly increased its workload in comparison with the one generated by the other disciplines. Generally speaking, during the first year of EGI-InSPIRE High-Energy Physics – and the ATLAS VO in particular (Figure 17) – was the main driver of resource utilization. This trend confirms what already observed since May 2009 (Figure 16 bottom).

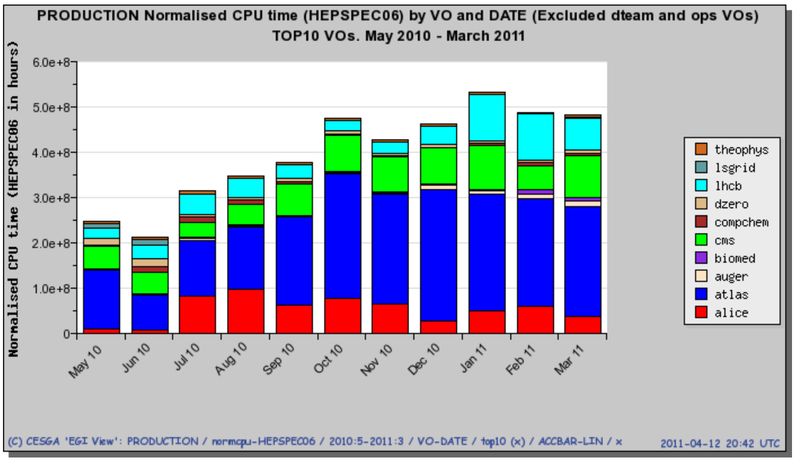

Figure 17. Distribution of consumed normalized CPU time (HESPEC06 hours) among the top ten VOs (source: Accounting Portal).

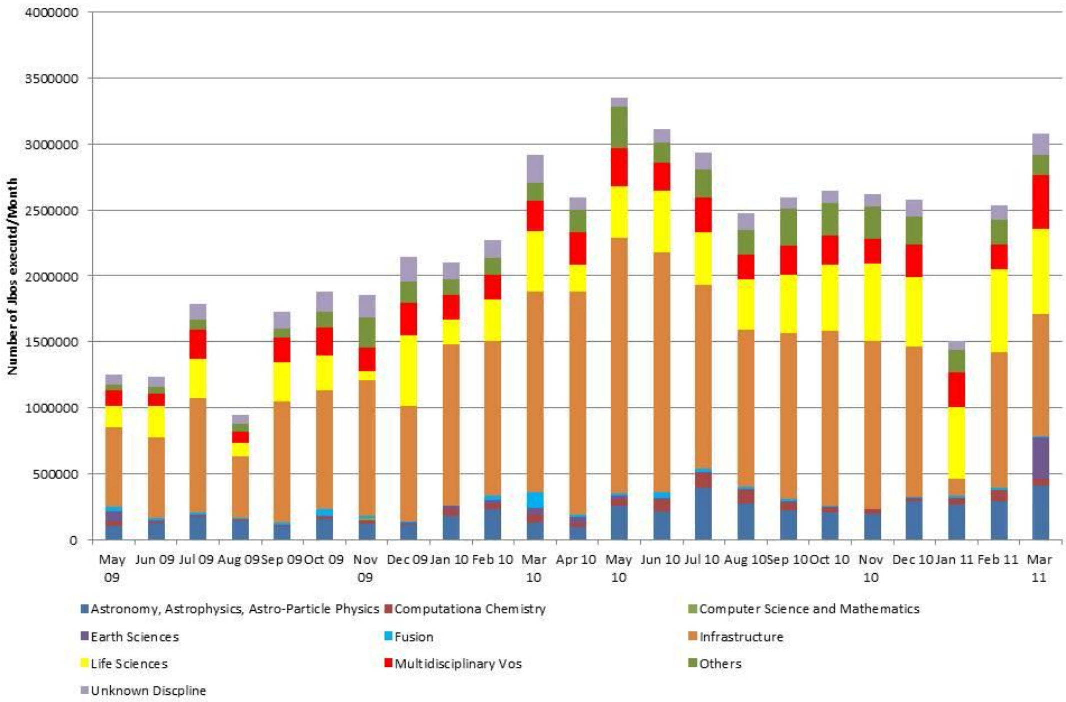

Figure 18. Number of executed jobs/month per non-LHC discipline from May 2009 (source: Accounting Portal)

The overall average number of jobs per month of non-HEP disciplines currently amounts to approximately 1.52 Mjobs/month (3.2 Mjobs/month if we just exclude LHC VOs) as shown in Figure 19). The overall number of jobs executed per month for non-HEP disciplines (Figure 18) since 2009 has approximately doubled from May 2009 to April 2010 (+100%). This trend continued in the following 11 months from May 2010 to March 2011, but at a much slower rate (+11%). The overall increase from May 2009 is +134%.

The following table summarizes the job statistics with and without HEP VOs, and compares the results with the averages from May 2009 to April 2011.

Table 7. Average number of jobs per day and per months with and without the HEP VOs

Metric	VOs	May 2010 - March 2011	Y1 target	May 2009 – April 2010
AVG number of job/day	All VOs	933,000	500,000	442,000
	No-HEP VOs	55,200	-	48,750
AVG number of Million job/month	All VOs	25.70	-	13.43
	No-HEP VOs	1.52		0.97

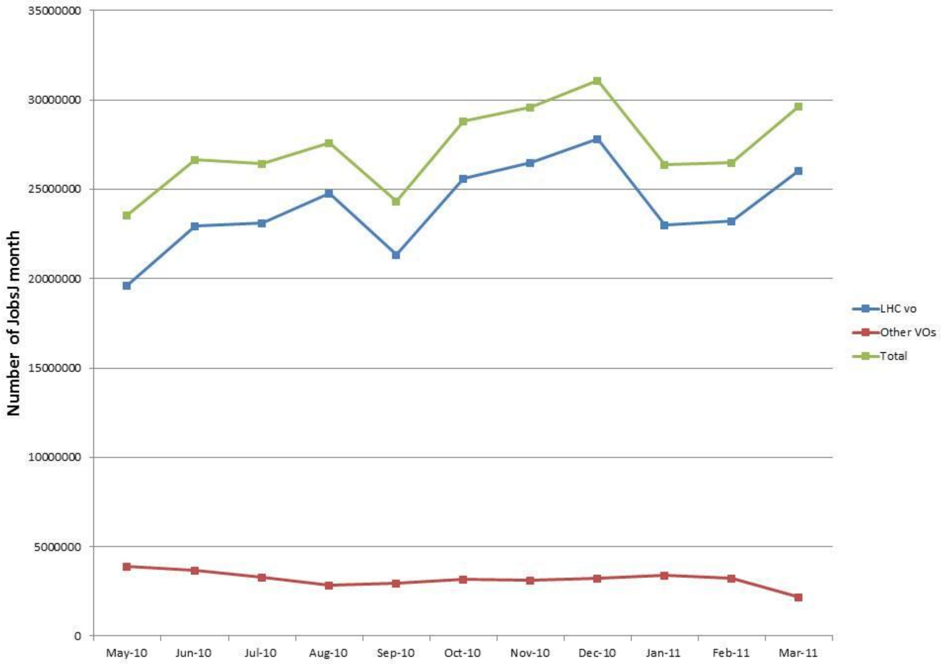

Figure 19. Comparison of jobs executed per month for LHC VOs and the cumulative workload of the other disciplines (source: Accounting Portal).

6 RESOURCE UTILIZATION

This section discusses the distribution of compute capacity utilization in EGI from May 2010.

Since the beginning of the project, the APEL-based accounting infrastructure has been under migration to adopt messaging (Apache ActiveMQ[17]) as communication channel to publish accounting data centrally. The messaging infrastructure is currently based on a geographically distributed network of message brokers operated at CERN, in Croatia and Greece. On the 1st of March 2011, the old central accounting repositories based on R-GMA, were decommissioned.

This migration only concerns the Resource Infrastructures based on APEL. Other infrastructures such as the Italian NGI, NGI_NDGF and OSG, where usage records are published through direct insertions into the central databases, and whose national accounting infrastructures adopt different technologies (e.g. DGAS, GRATIA and SGAS), are not involved in this transition.

Currently 80% of the R-GMA Resource Centres have completed the migrated to messaging. On the 1st of March the central R-GMA databases were decommissioned, and since then Resource Centres where the APEL accounting client has not been updated, cannot publish accounting data centrally. In these cases accounting records are accumulated locally and ready for being published as soon as the R-GMA client will be upgraded. The decommissioning of the central databases and the ongoing transition are reflected in the accounting statistics from March 2011 (included). These are incomplete and the amount of missing usage records is proportional to the number the Resource Centres still to be migrated, and their usage.

The accounting statistics from March 2011 onwards will be gradually rectified as more Resource Centres will migrate to the ActiveMQ APEL client.

6.1 Infrastructure Workload Distribution

The total consumed capacity from May 2010 to March 2011 amounts to 4.53 Billion HEP-SPEC06 Hours.

EGI accounting information is gathered and stored centrally for display through the accounting portal. Accounting information is aggregated by Operations Centre, whose list is obtained from GOCDB. The distribution of consumed compute capacity – in Million HEP-SPEC06 Hours – from May 2010 to March 2011 is illustrated in Figure 20, where the utilization per Operations Centre is displayed.

For Operations Centres that became operational after May 2010, the accounting statistics herein presented are complete, as they include all the accounting information available for the respective Resource Centres starting from May 2010, even if for part of the time centres were affiliated to a different entity. This is because the accounting information per Operations Centre is currently computed by aggregating the entire set of accounting records from Resource Centres, even if for a part of the reference time interval the Resource Centres were affiliated to a different Operations Centre.

[17] http://activemq.apache.org/

Figure 21 (a) compares the overall compute resource utilization in HEP-SPEC06 Hour from May 2010 to date for the active EGI Operations Centres. The resource infrastructures that score more than 10% of utilization are those operated by the UKI Federation (16.31%), NGI_FR (15.47%), NGI_DE (14.23%) and Italy (12.66%)[18]. Figure 21 (b) compares the utilization per country including United States of America (OSG). The normalized CPU time utilization percentage of the integrated EGI-InSPIRE infrastructure (including integrated Resource Infrastructure Providers that are consumers of EGI services) is 70.8% of the total, while the OSG utilization amounts to 29.2%.

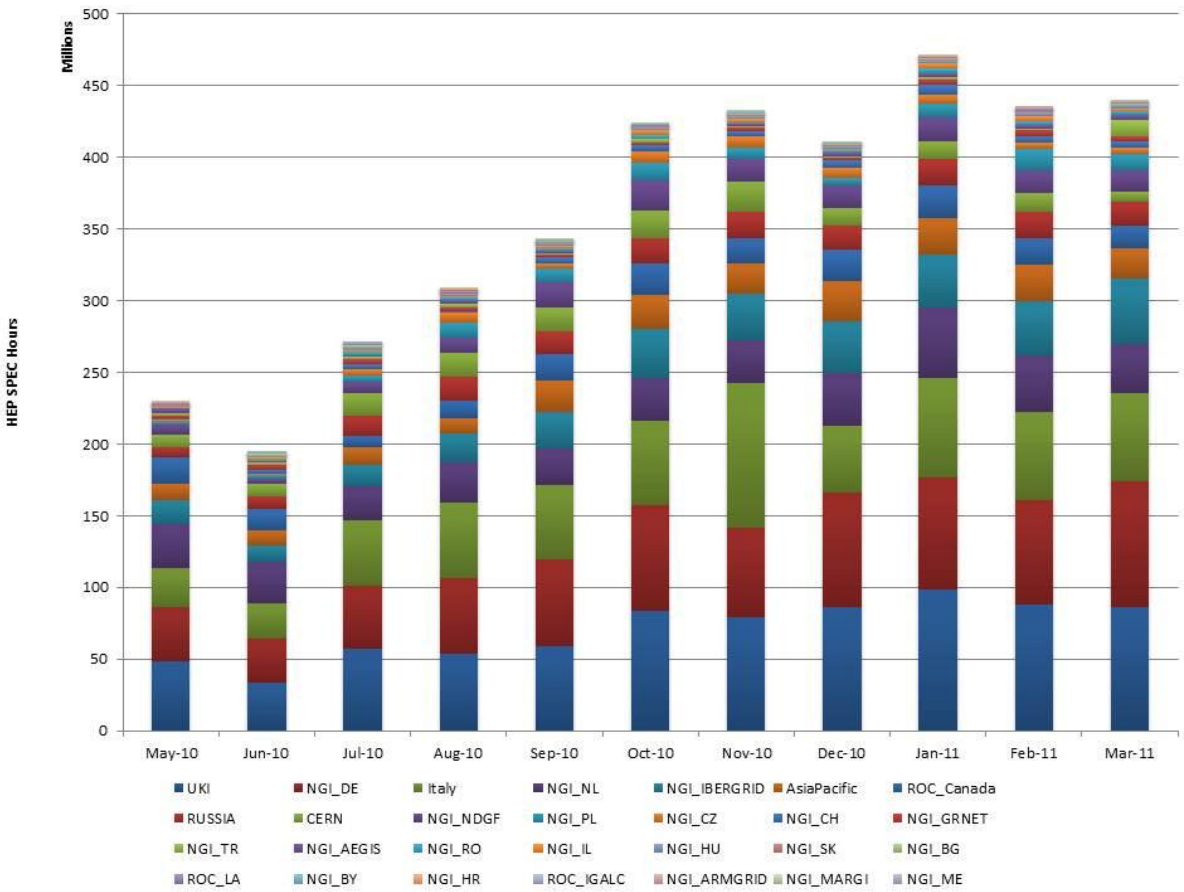

Figure 20. Distribution of consumed compute capacity (Millions of HEP-SPEC06 Hours) between the EGI Operations Centres from May 2010 to date (source: Accounting Portal)

---

[18] CERN accounting statistics from November 2010 are known to be underestimated by a factor of 2.5 due to a wrong normalization factor. The problem is being worked upon.

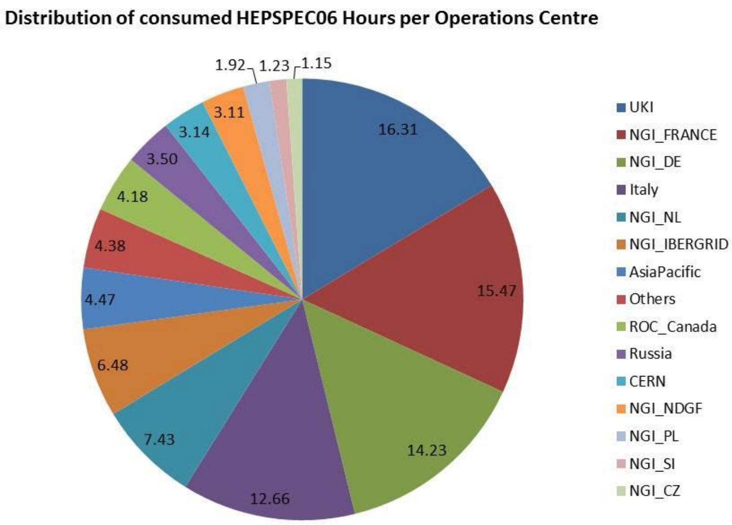

(a)

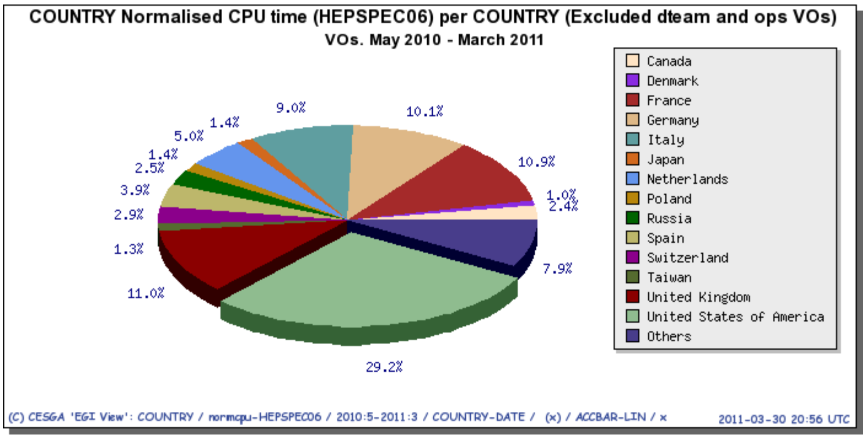

(b)

Figure 21. Percentage of compute capacity utilization (HEP-SPEC06 Hours) per EGI Operations Centre (a) and compute capacity utilization (HEP-SPEC06 Hours) per country including peer Resource Infrastructures (b)- from May 2011 to date. The "Others" category in diagram (a) includes all infrastructures whose utilization was less then 1%.

6.2 VO Workload Distribution

The distribution of VO workload among the various Resource Infrastructures varies considerably depending on the amount of resources offered locally and by the distribution of the end-users.

As shown in Table 8, several infrastructures (or federations of infrastructures) are heavily driven by the LHC community, where LHC resource usage varies in the range [90, 100] %: the Canada Federation, Romania, CERN, Switzerland, Israel, the NDGF Federation, Russia, Poland, Slovakia, Slovenia and the

IberGrid Federation. Others show a more balanced composition with a LHC utilization ranging in the interval [75, 90]%, these are in descending order: Germany, the United Kingdom/Ireland Federation, Czech Republic, Hungary, France, the Asia Pacific federation, the Latin America Federation and Croatia. Finally, the remaining infrastructures where the LHC utilization of resources is less that 75%, and hence more balanced in comparison with other disciplines, are in descending order: Italy, Bulgaria, Netherlands, the GISELA Federation, Cyprus, Turkey, Serbia, Greece, Belarus, Armenia, Montenegro and FYR of Macedonia.

Table 8. LHC workload per NGI/EIRO/Federation (normalized CPU time) from May 2010 to Mar 2011 in descending order.

Operations Centres	LHC utilization	Operations Centres	LHC utilization	Operations Centres	LHC utilization
ROC_Canada	99.5%	NGI_DE	89.7%	ROC_IGALC	67.2%
NGI_RO	98.5%	UKI	87.8%	NGI_CYGRID	36.7%
CERN	98.1%	NGI_CZ	87.7%	NGI_TR	33.4%
NGI_CH	97.9%	NGI_HU	83.7%	NGI_AEGIS	15.2%
NGI_IL	97.4%	NGI_FRANCE	82.9%	NGI_GRNET	9.8%
NGI_NDGF	97.3%	Asia Pacific	81.1%	NGI_BY	4.3%
Russia	95.8%	ROC_LA	77.4%	NGI_ARMGRID	0.03%
NGI_PL	93.5%	NGI_HR	76.7%	NGI_ME	0.00%
NGI_SK	93.5%	Italy	74.8%	NGI_MARGI	0.00%
NGI_SI	92.6%	NGI_BG	72.4%		
NGI_IBERGRID	91.0%	NGI_NL	71.8%		

The detailed usage breakdown per VO for a sample of the EGI Infrastructures is illustrated in Figure 22 below.

Figure 22. Detailed sample of workload distribution per NGI/Federation injected by the top VOs (normalized CPU time) from May 2010 to March 2011. The Infrastructures included are those with an overall utilization larger than 3% (May 2010 – May 2011) – source: Accounting Portal.

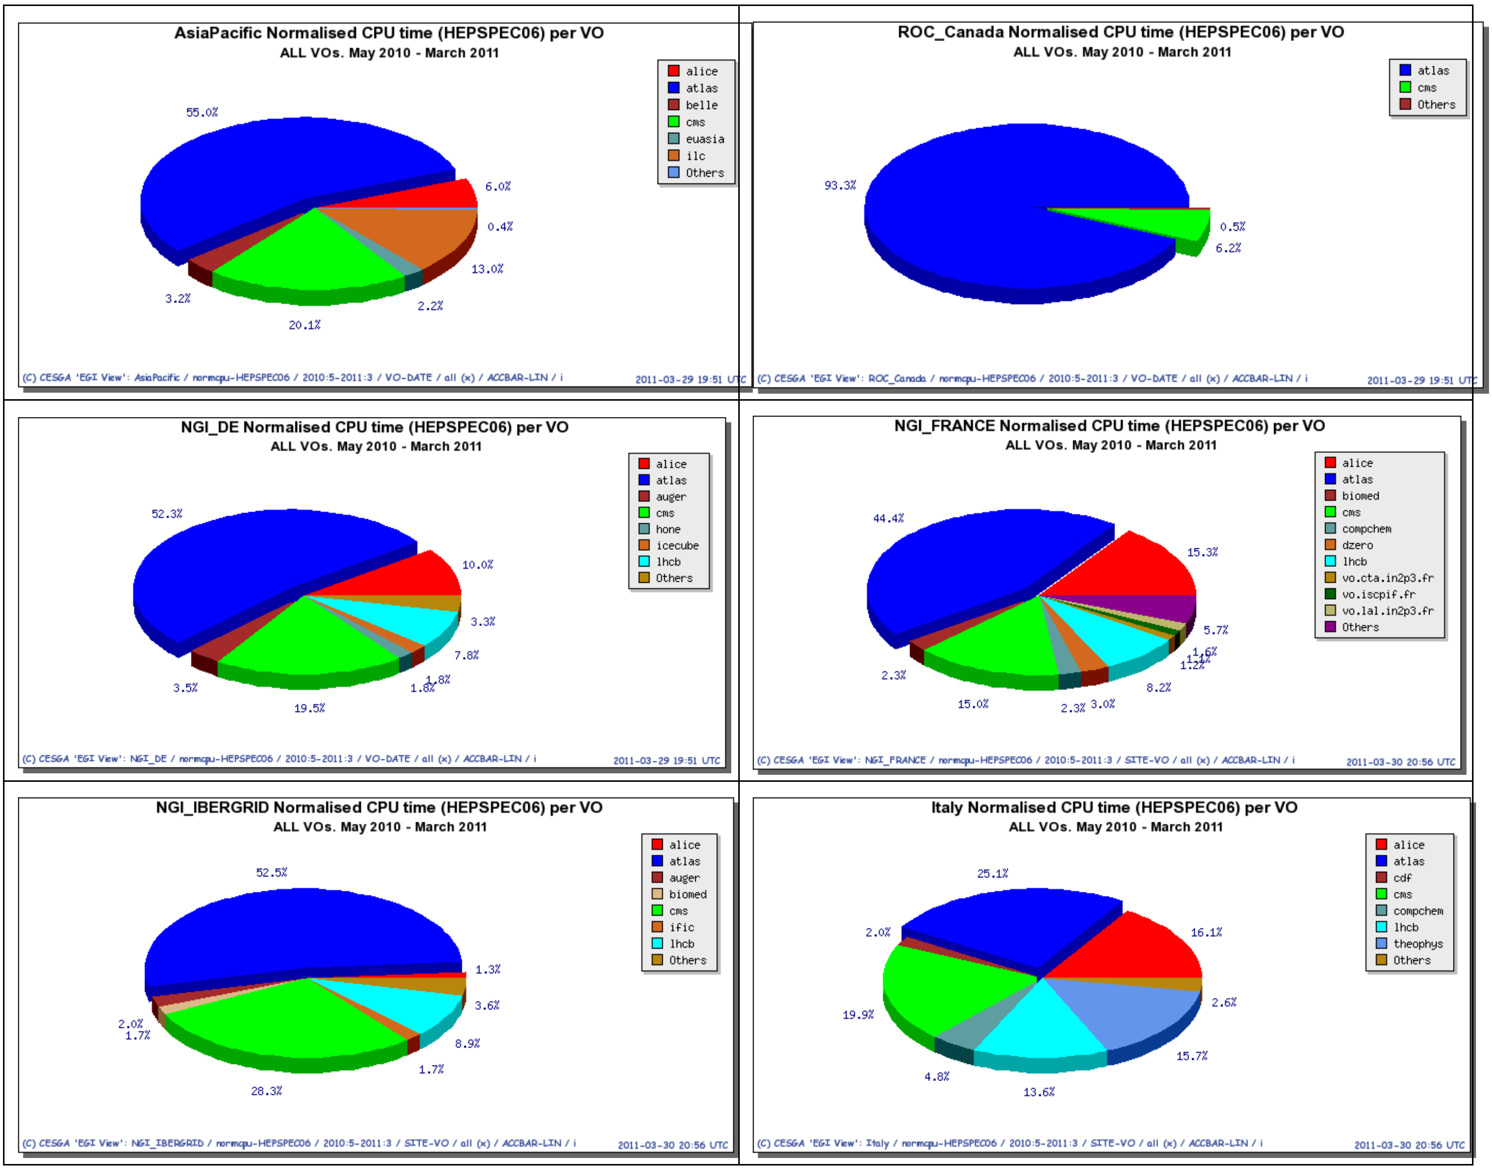

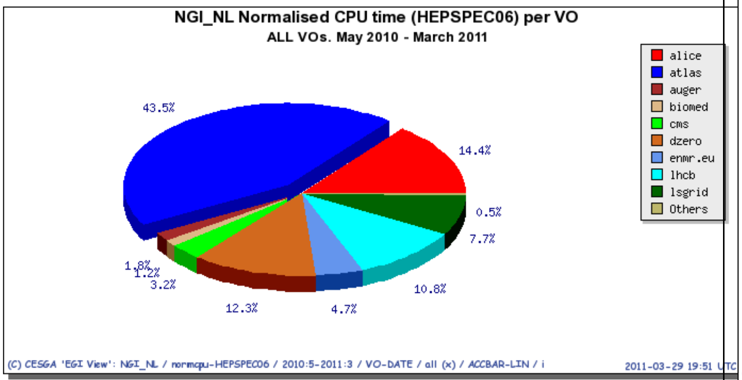
NGI_NL Normalised CPU time (HEPSPEC06) per VO
ALL VOs. May 2010 - March 2011
alice
atlas
auger
biomed
cms
dzero
enmr.eu
lhcb
lsgrid
Others
43.5%
14.4%
0.5%
7.7%
10.8%
4.7%
12.3%
3.2%
1.8%
1.2%
(C) CESGA 'EGI View': NGI_NL / normcpu-HEPSPEC06 / 2010:5-2011:3 / VO-DATE / all (x) / ACCBAR-LIN / i
2011-03-29 19:51 UTC

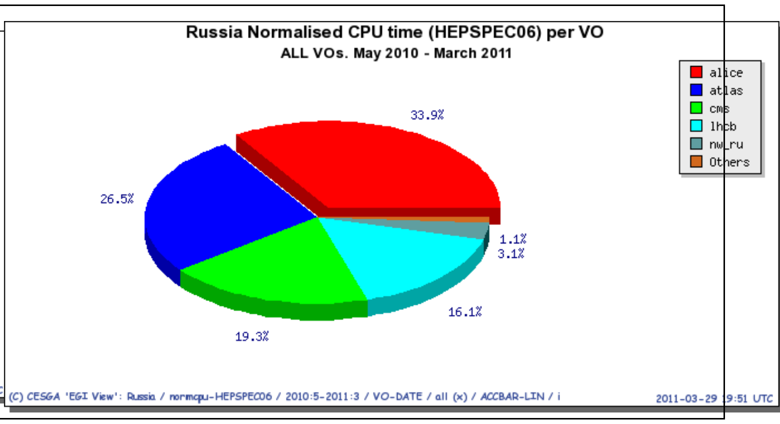
Russia Normalised CPU time (HEPSPEC06) per VO
ALL VOs. May 2010 - March 2011
alice
atlas
cms
lhcb
nw_ru
Others
33.9%
26.5%
1.1%
3.1%
16.1%
19.3%
(C) CESGA 'EGI View': Russia / normcpu-HEPSPEC06 / 2010:5-2011:3 / VO-DATE / all (x) / ACCBAR-LIN / i
2011-03-29 19:51 UTC

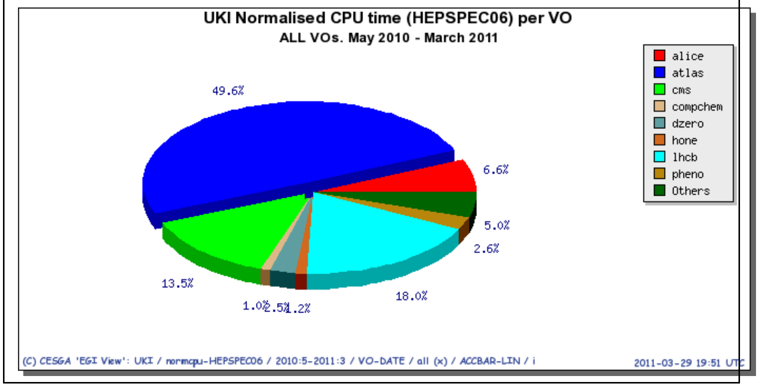
UKI Normalised CPU time (HEPSPEC06) per VO
ALL VOs. May 2010 - March 2011
alice
atlas
cms
compchem
dzero
hone
lhcb
pheno
Others
49.6%
6.6%
5.0%
2.6%
18.0%
13.5%
1.0%
2.5%
1.2%
(C) CESGA 'EGI View': UKI / normcpu-HEPSPEC06 / 2010:5-2011:3 / VO-DATE / all (x) / ACCBAR-LIN / i
2011-03-29 19:51 UTC

7 DEPLOYED MIDDLEWARE

7.1 Middleware Stacks

One of the objectives of EGI-InSPIRE is “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” (Objective 5). In the first project year the SA1 focus has been the integration of existing European infrastructures based on different middleware stacks with the operational infrastructures: accounting, monitoring, management and support. Activities have been concentrated on the integration of ARC and Globus resources into the Nagios[19]-based monitoring infrastructure, and on the integration of UNICORE resources into the EGI central registry (GOCDB) and monitoring.

While the deployment of gLite and ARC were already consolidated during the EGEE project series, the integration of UNICORE and Globus resources are at their infancy. As shown in Figure 23, in March 2011 the gLite-based Operations Centres are 37. ARC and UNICORE are supported and deployed in one Operations Centre: NGI_NDGF and NGI_DE.

NGI_NDGF includes Denmark, Finland, Norway and Sweden (all ARC-based) and additional countries currently deploy ARC test instances, these are Baltic countries, Slovenia and Switzerland. NGI_NDGF currently deploys and support two middleware stacks: ARC and gLite.

Several other countries are planning to integrate and support multiple middleware stacks, such as: NGI_DE (gLite, Globus and UNICORE), NGI_NL (gLite and Globus), NGI_PL (gLite, Globus and UNICORE), NGI_RO (gLite, Globus, UNICORE), IberGrid (gLite and Globus), and United Kingdom (gLite and Globus). Additional Resource Infrastructure Providers have expressed interest in supporting more middleware stacks in case of demand.

[19] http://www.nagios.org/

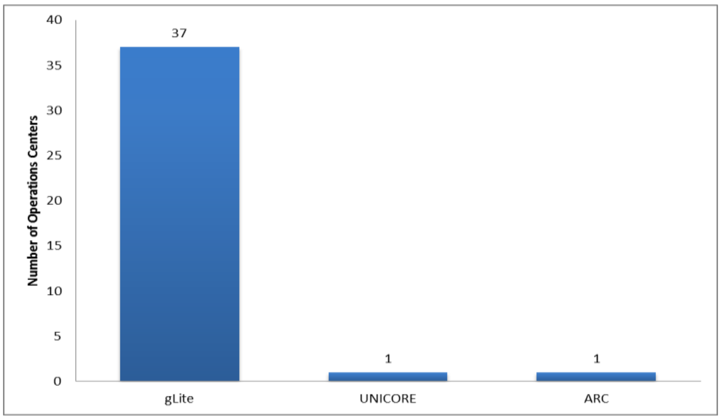

Figure 23. Deployment of the four reference grid middleware stacks (gLite, ARC and UNICORE) across the EGI-InSPIRE Operations Centres – March 2011 (source: GridMap[20]). Note that federated Operations Centres provide services to multiple Resource Infrastructures.

Figure 24. Distribution of different implementations of the Compute Capability (ARC CE, CREAM CE and lcg-CE) across the EGI-InSPIRE partners and the integrated Resource Infrastructure Providers – March 2011 (source: GridMap).

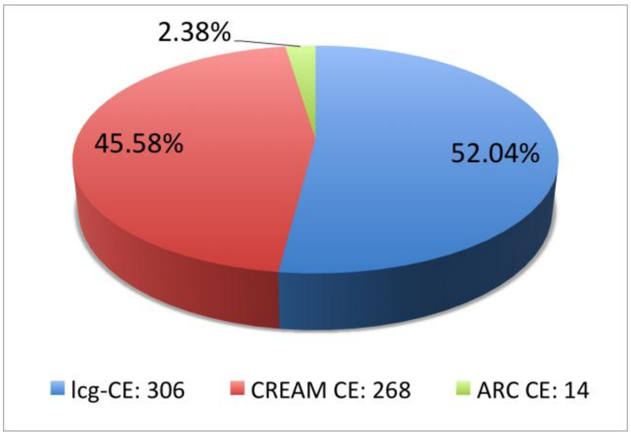

Figure 24. Distribution of different implementations of the Compute Capability (ARC CE, CREAM CE and lcg-CE) across the EGI-InSPIRE partners and the integrated Resource Infrastructure Providers – March 2011 (source: GridMap).

The level of integration of the different implementations of the compute capability is shown in Figure 24, which compares the amount of production instances of ARC CE, CREAM CE and lcg-CE.

lcg-CE is instances currently constitute the majority, they amount to 52.04 % of the existing Compute Elements. As end-of-support of lcg-CE is foreseen to be scheduled by the end of 2011, this will require the Resource Infrastructure Providers to gradually migrate to different implementations.

[20] http://gridmap.cern.ch/

7.2 Core Middleware Services

Resource Infrastructure Providers are required to operate core Grid middleware services according to the needs of the local and global VOs that are supported locally. In this context the term *core middleware service* refers to a Grid technical service providing capabilities necessary to access resources at the Resource Centre-level. Core services can serve multiple VOs or are dedicated, depending on the workload generated by the end-users. Examples of such services are file catalogues (e.g. gLite LFC), user authentication (e.g. gLite VOMS), the information discovery System (e.g. gLite top-BDII) and workload management (e.g. gLite WMS). Other Grid services provide core capabilities, however in this section we restrict our analysis to four of them.

This section provides a snapshot of the current distribution of production instances of LFC, top-BDII, VOMS and WMS services deployed across EGI (including both EGI-InSPIRE partners and integrated Resource Infrastructures). Information was extracted from the top-level Information Discovery System. The overall number of EGI core services amounts to 339 units. 135 of them are WMS instances (this count comprises WMS services needed by the Nagios-based monitoring infrastructure), 45 are gLite LFC file catalogues, 118 are top-BDII instances and 41 are VOMS servers. Table 9 shows the distribution of those services per country.

Table 9. Number of core middleware service instances deployed across EGI-InSPIRE partners and the integrated Resource Infrastructures – March 2011 (source: top-BDII).

(a) **Workload Management** gLite WMS	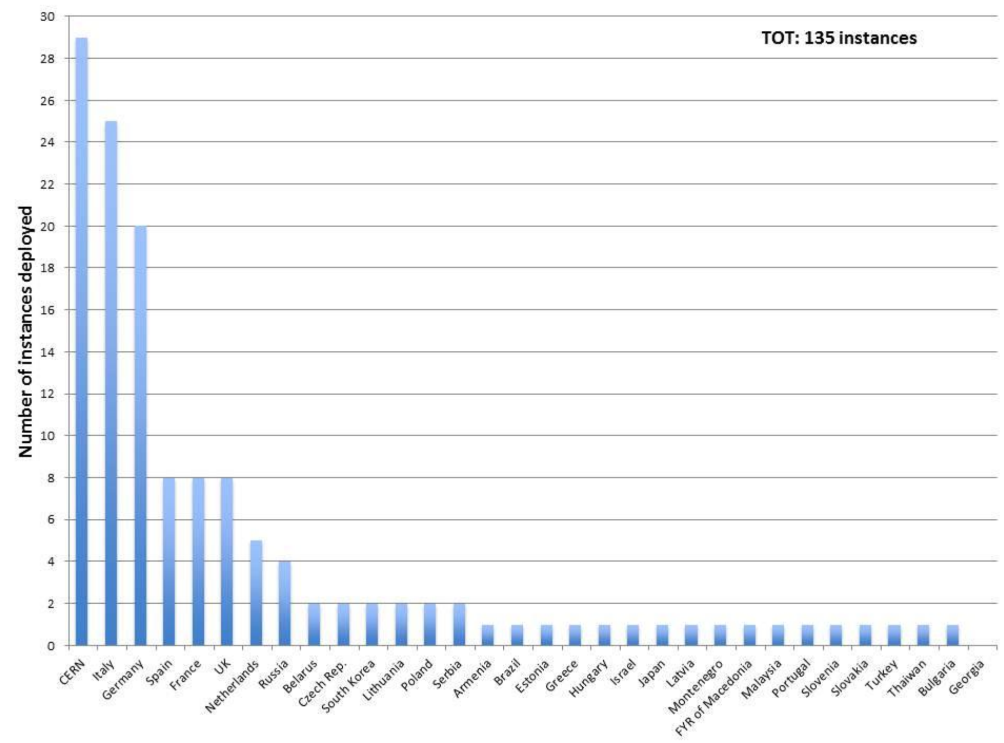

<table>
<tr>
<td>(b)
File catalogues
gLite LFC</td>
<td>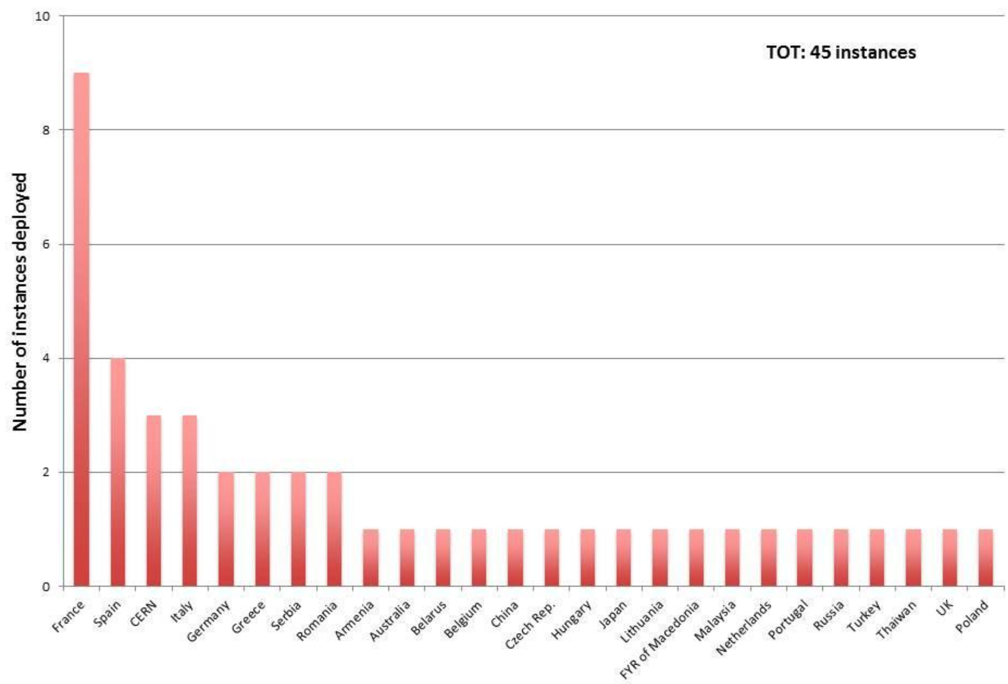
</td>
</tr>
<tr>
<td>(c)
Information Discovery
Glite Top-BDII</td>
<td>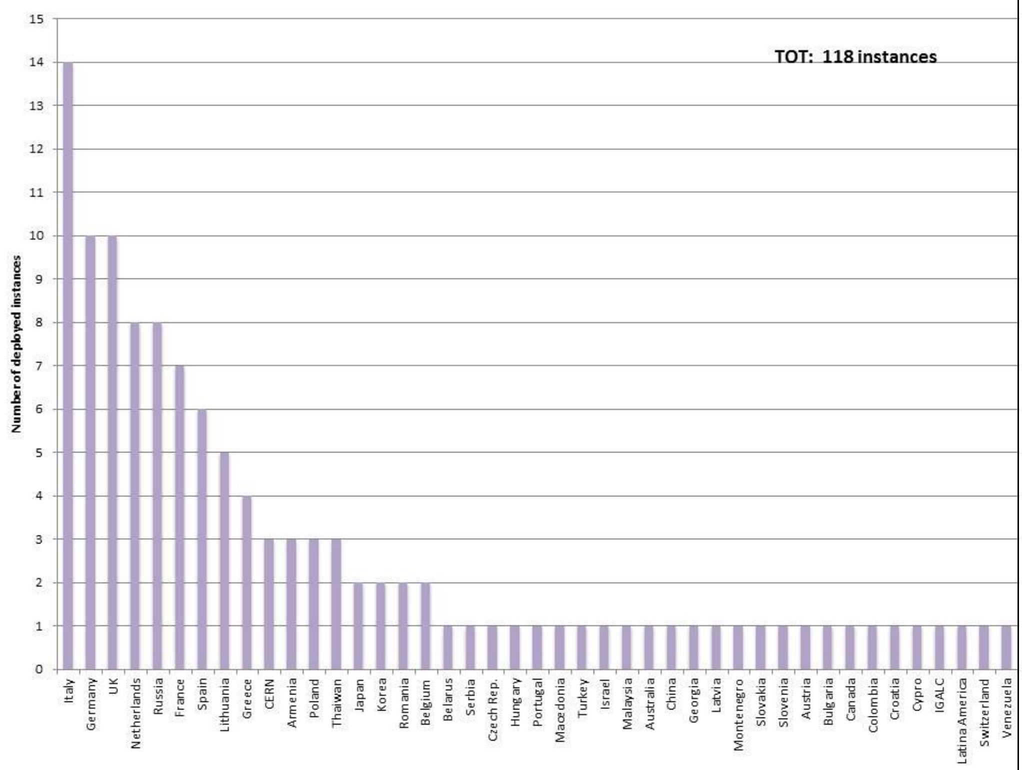
</td>
</tr>
</table>

<table>
<tr>
<td>(d)
VO Membership
gLite VOMS</td>
<td>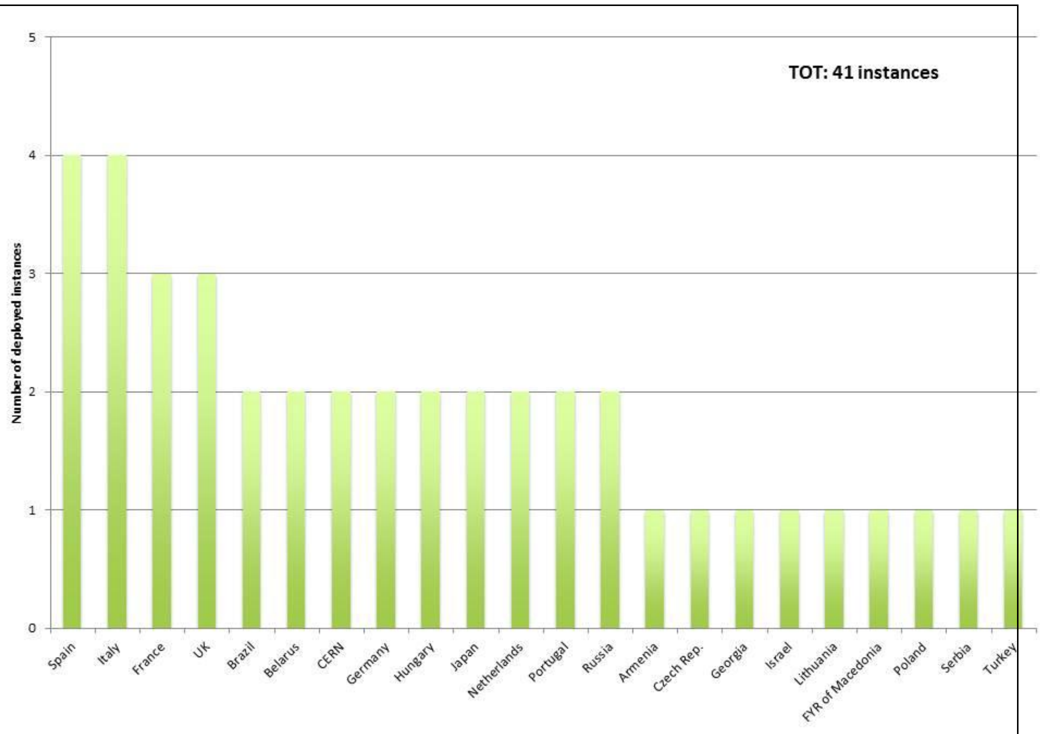
</td>
</tr>
<tr>
<td>(e)
Overall distribution</td>
<td>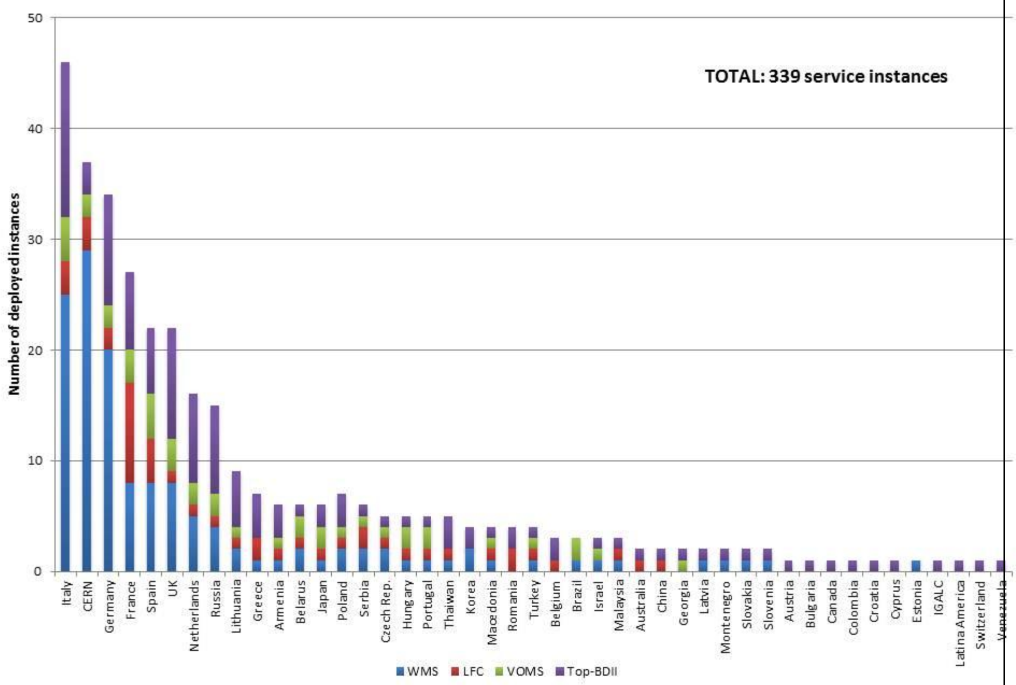
</td>
</tr>
</table>

(f) **Overall distribution and number of active VOs** (source of active VOs: Accounting Portal)	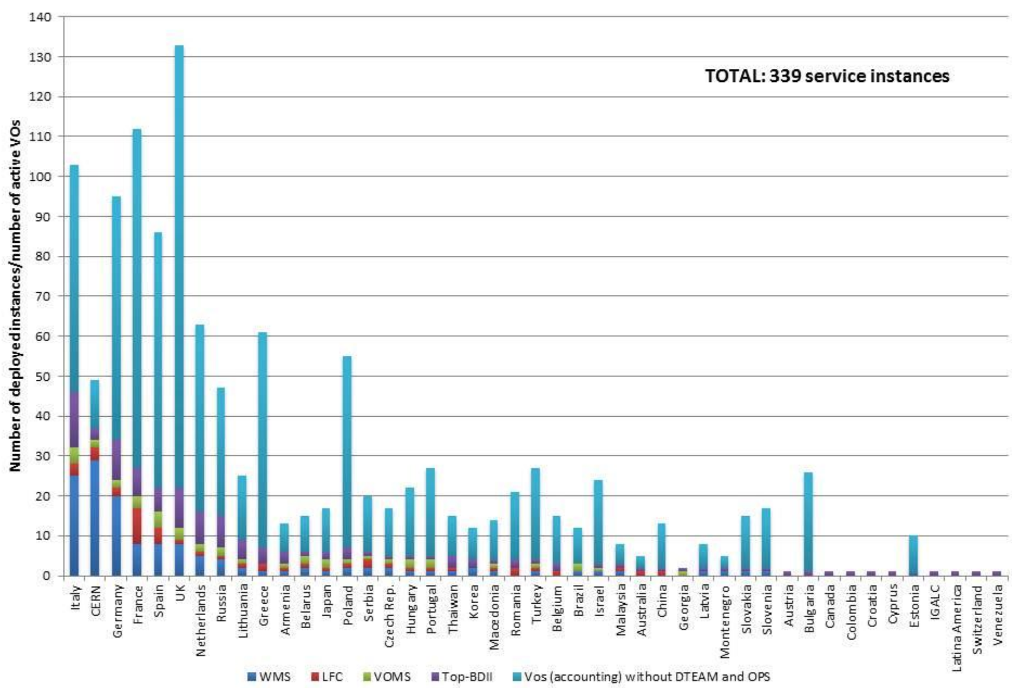

As shown in Table 9 (e), the Resource Infrastructure with the largest amount of services is Italy (46 instances), followed by CERN (37 instances), Germany (34), France (27), Spain (22), and United Kingdom (22) followed by the other countries in the range [1, 20] instances. As expected, the number of services is correlated with the amount of installed capacity provided and with the size and activity of the supported VOs. For larger communities in many cases dedicated entries of core middleware services are needed. In addition, being these core services, high availability needs to be provided to be provided to the users. To this end, many of the core services are deployed in cluster mode to improve resiliency to fault of individual instances and load balancing.

Table 9 (f) shows the correlation between the total number of middleware services operated and the number of active VOs, where with active we mean VOs for which accounting information is available. CERN is an example of Resource Infrastructure that includes a large number of core services for its heavy users, but which actually support a relatively limited number of communities. Note that not all of the VOs included in this count actually generated accounting from March 2010, this meaning that the number does not reflect the amount of currently active VOs.

7.3 Compute and Storage Management Services

Several software components from the gLite 3.1 middleware distribution are currently reaching end of standard and/or security support [CAL], this has operational implications as a plan of upgrade is being put in place by the Resource Infrastructure Providers. Due the upcoming release of EMI 1.0 (end of April 2011), an upgrade from gLite 3.1 to an equivalent EMI 1.0 component seems to be the best option, as a replacement of a gLite 3.1 service with a gLite 3.2 version (where available) requires a change of the Operating System.

As shown in Table 10, nearly all worker nodes operated by the Resource Centres has already migrated to gLite 3.2 (a). The situation is considerably different for stateful services such as LFC (b) and DPM (c) where the percentage of gLite 3.1 components amount to 38.64 % and 39.65 % respectively. As to WMS, a campaign is in progress in order to accelerate the update of instances to version 3.1.30; this upgrade is necessary as the new release includes a new version of gridsite, which fixes a problem happening when considering proxies generated using gLite 3.2 VOMS servers[21].

Table 10. Comparison of deployment of gLite 3.1 services and gLite 3.2 services: WN (a), LFC (b), DPM (c) and WMS (d) – March 2011 (source: top-BDII).

<table>
<tr><td>gLite-WN
Only Security Updates
until 30/04/2011</td><td>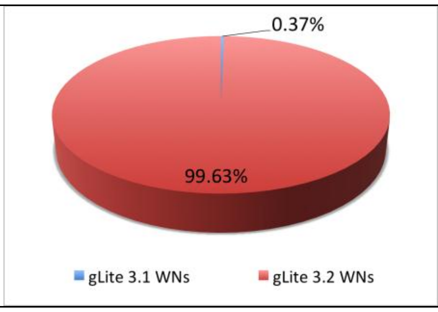
</td></tr>
<tr><td>(b)
gLite- LFC
End of Standard Updates
on 09/04/2011</td><td>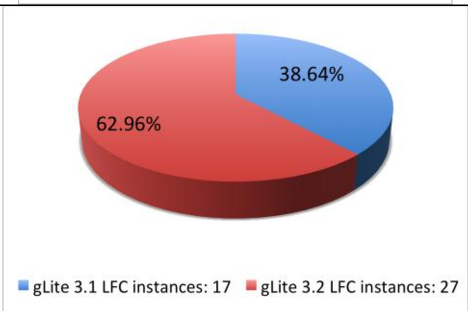
</td></tr>
<tr><td>(c)
gLite-DPM
Only Security Updates
until 18/05/2011</td><td>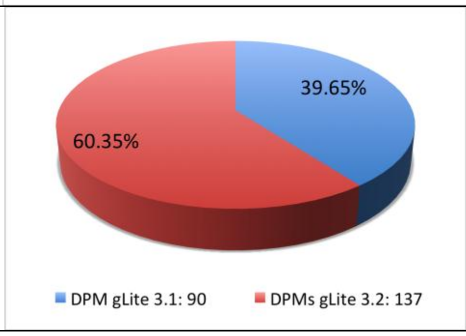
</td></tr>
</table>

[21] http://go.egi.eu/wms3.1.30

(d) gLite-WMS (gLite 3.1 only), necessary gridsite fix in version 3.1.20	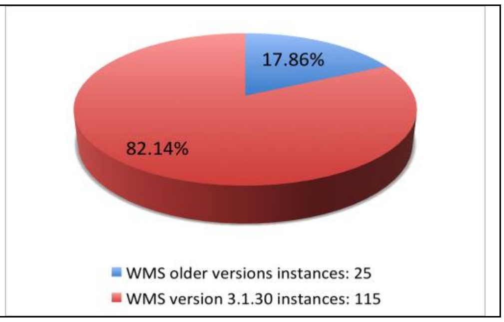

8 STAGED ROLLOUT INFRASTRUCTURE

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 has been introduced and 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.

Table 11. Overview of EGI-InSPIRE Staged Rollout activities (PQ1, PQ2 and PQ3).

Metric	PQ1	PQ2	PQ3
Total number of patches ready for Staged Rollout	18	17	11
Patches tested by one or more EAs	11	10	11
Patches rejected	1	3	2
Patches rolled to production without Staged Rollout (fixing a critical incident, or without EAs)	7	7	0
Number of staged rollout tests undertaken	12	14	14

Table 11. Overview of EGI-InSPIRE Staged Rollout activities shows the staged rollout activities in the first EGI year. In the first two PQs not all the components entering the Staged Rollout stage were actually validated, either because of the lack of EAs or the urgency of the fix introduced with the patch (one case). This has been gradually improving. During PQ3 all the patches in the Stage Rollout

have been properly processed by the corresponding EAs, and currently the goal to have full coverage for every software component deployed in EGI.

The Staged Rollout Infrastructure currently comprises 39 Resource Centres. Several EAs are participating to the Staged Rollout of two or more software components. Presently the Staged Rollout Infrastructure covers completely the gLite (release 3.1 and 3.2) and ARC middleware stacks, most of UNICORE and Globus components, and one Operational Tool (the SAM framework and Nagios probes[22]).

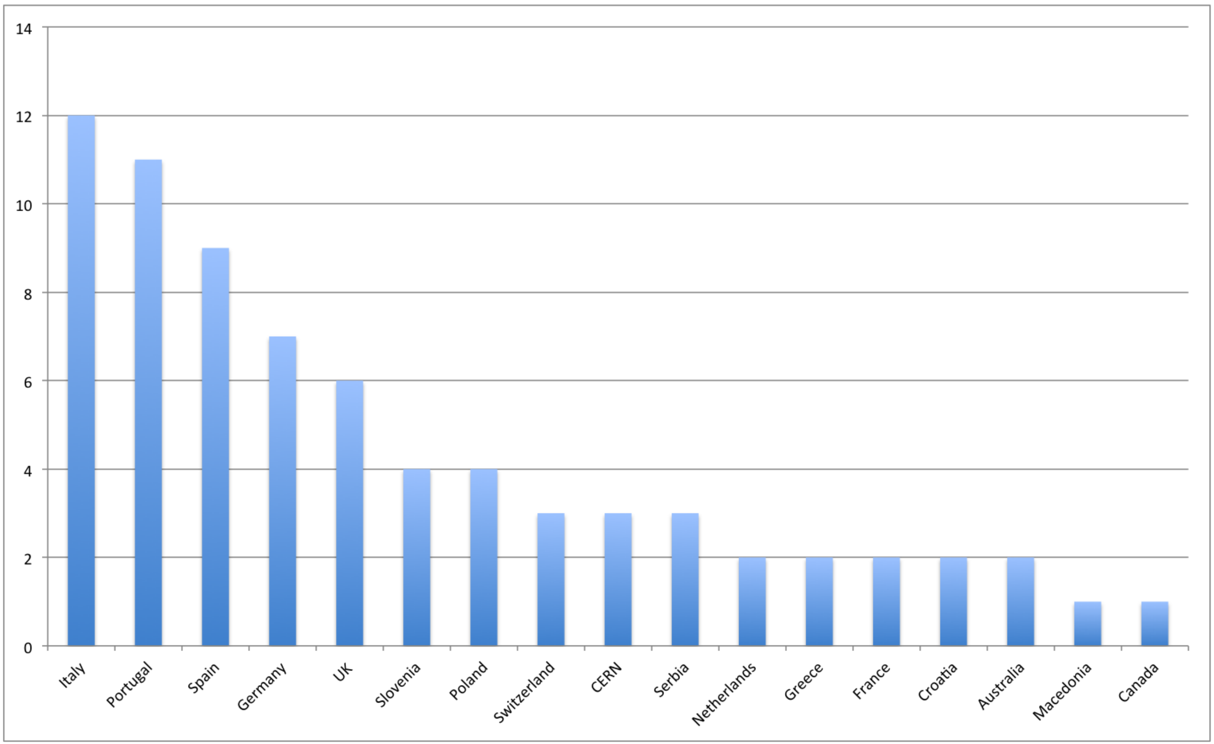

Figure 25. Number of staged rollout tests by NGI (May 2010 - March 2011).

Figure 25 shows the number of staged rollout tests in the last year operated by the NGIs. Having multiple EAs for a given software is important as this increases the chance to expose it to different configuration, deployment and usage scenarios. This contributes to increase the effectiveness of the overall process. The number of EAs that participate in the Staged testing of a patch currently varies between 1 and 9.

[22] https://wiki.egi.eu/wiki/SAM

9 CONCLUSIONS AND FUTURE PLANS

In March 2011 the Resource Infrastructure status satisfactorily met all the infrastructure targets defined for the first year of EGI-InSPIRE as to the number of Resource Centres integrated and the number of those offering high-performance computing and MPI capabilities, the amount of logical CPUs contributed, the average monthly reliability, and utilization (average number of jobs per day).

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

  This objective was successfully met. Performance was granted during a radical transition process that evolved the 14 federated Operations Centres existing in April 2010 into 40 NGIs operated by 27 Operations Centres. During this transition the installed capacity and Resource Centres integrated continued to grow. During year 2011 effort will concentrate on the automation of operational tools and services for the improvement of the infrastructure performance.

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

  Progress was made to extend the number of integrated middleware stacks. In March 2011 gLite and ARC are fully integrated. The feasibility of the integration of UNICORE and Globus with the existing accounting, monitoring and tooling infrastructures was demonstrated. One UNICORE Resource Centres is now fully part of the production infrastructure. In year 2011 effort will be focused on the full integration of UNICORE and Globus. The integration of new infrastructures is expected to facilitate and foster the support of new user communities for an increasingly balanced resource usage by different scientific disciplines.

- 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*.

  EGI-InSPIRE is currently collaborating with three non-European Resource Infrastructure Providers for the definition of a "Resource Infrastructure Provider MoU" to facilitate the exchange of operational services and the integration between the EGI-InSPIRE infrastructure and those operated by external partners. 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 number of integrated high-performance cluster is also increasing. 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. The activity participates in the Technology Coordination Board to present requirements to the Technology Providers for the seamless integration of multiple software stacks that can further improve uptake by users and Resource Centres. In additions discussions are ongoing to understand how to integrate virtualized resources and desktop Grids.

10 REFERENCES

D4.1	EGI Operations Architecture, EGI-InSPIRE Deliverable D4.1, Jan 2011 (https://documents.egi.eu/secure/ShowDocument?docid=218)
MS108	Global Task Review, EGI-InSPIRE Milestone MS108, Mar 2011 (https://documents.egi.eu/secure/ShowDocument?docid=314)
MS109	NGI International Task Review, EGI-InSPIRE Milestone MS109, Mar 2011 ()
MS407	Integrating Resources into the EGI Production Infrastructure, EGI-InSPIRE Milestone MS407, Nov 2010 (https://documents.egi.eu/document/111)
UMD	UMD Roadmap, EGI-InSPIRE Deliverable D5.1, Aug 2010 (https://documents.egi.eu/document/100)
MOU	Resource Infrastructure Provider MoU template (https://documents.egi.eu/secure/ShowDocument?docid=215)
SOP	EGI Grid Site Operations Policy, Jul 2010 (https://documents.egi.eu/document/75)
AVA	Availability and Reliability monthly Statistics: (https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics)
PROC04	Process for Quality Verification, EGI Operations Procedure PROC04, Jul 2010 (https://wiki.egi.eu/wiki/Availability_and_reliability_monthly_statistics)
SEC03	EGI-CSIRT Critical Vulnerability Handling, EGI Security Procedure SEC03, Mar 2011 (https://wiki.egi.eu/wiki/SEC03)
MAN02	Service Intervention Management, EGI Operations Manual MAN02 (https://wiki.egi.eu/wiki/MAN02), Mar 2011
MPI	MPI: A Message-Passing Interface Standard Version 2.2, Message Passing Interface Forum, Sep 2009 (http://www.mpi-forum.org/docs/mpi-2.2/mpi22-report.pdf)
HS06	The HEPiX CPU Benchmark Working Group Wiki (http://hepix.caspur.it/processors/dokuwiki/doku.php?id=benchmarks:introduction)
HWG	The HEPiX CPU Benchmark Working Group Wiki (http://hepix.caspur.it/processors/dokuwiki/doku.php?id=benchmarks:introduction)
SRW	Staged Rollout Workflow (https://wiki.egi.eu/wiki/Staged-rollout-procedures)

GOC	Mathieu, G., Casson, J. (2010) GOCDB4, a New Architecture for the European Grid Infrastructure; in the Proc. of the Grid Computing Int. Symposium on Grid Computing (ISGC 2010), Taipei, Taiwan
CAL	The gLite distribution Support Calendar (http://glite.cern.ch/support_calendar/)