EGI-Engage

Report on the evolution of the EGI Operations Infrastructure

D5.1

Date	10 March 2016
Activity	SA1
Lead Partner	EGI.eu
Document Status	FINAL
Document Link	https://documents.egi.eu/document/2670

Abstract

This document presents the status of the EGI Operations infrastructure at the end of the first year of EGI-Engage, and reports on the improvements and evolution of services and processes that has been implemented during the first reporting period. This deliverable gives an overview of the deployed services that are provided to users and members of the federation, the amount of resources available, trends, and the activities that support service provisioning across the EGI federation.

The roadmap for evolving the current set of services is presented as well with links to other EGI-engage work packages.

The EGI-Engage project is co-funded by the European Union (EU) Horizon 2020 program under Grant number 654142 http://go.egi.eu/eng

COPYRIGHT NOTICE

DELIVERY SLIP

	Name	***Partner/Activity***	***Date***
From:	Peter Solagna	EGI.eu/SA1	9/03/2016
Moderated by:	Małgorzata Krakowian	EGI.eu/NA1	
Reviewed by	Kostas Koumantaros Tiziana Ferrari	GRNET/PMB EGI.eu/NA1	27/02/2016 1/03/2016
Approved by:	AMB and PMB		9/03/2016

DOCUMENT LOG

Issue	***Date***	***Comment***	***Author/Partner***
v.1	19/02/2016	First version	Peter Solagna/EGI.eu
FINAL	9/03/2016	Version after external review	Tiziana Ferrari/EGI.eu
	10/03/2016	New version of document produced with corrected accounting figures for the EGI Federated Cloud (section 2.6)	Tiziana Ferrari/EGI.eu

TERMINOLOGY

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

Contents

Executive summary

EGI is the largest international e-infrastructure providing advanced compute, storage, data management and software and service platforms that provide scientists from any discipline with the digital services needed for research and innovation in Europe and worldwide. In the first quarter of 2016, the capacity deployed in EGI exceeded the thresholds of 650,000 logical CPUs (23.6% yearly relative increase compared to 8.13% of the previous year) and 500 PB of online and near-line storage federated worldwide across 58 countries, more than 325 data centres in Africa, Asia-Pacific region, Canada, Europe and South America. 343200 Virtual Machines (+79% yearly increase, 2.31 Million hours of CPU wall time consumed) were instantiated in the Federated Cloud from Jan 2015 to Jan 2016. The federated cloud, with its 21 providers of which one is commercial, is a new platform that started its production activities in May 2014.

In 2016 Q1 the estimated number of active users exceeded 57,000, of which 66% active in natural sciences; 6.5% in medical and health research and 6.3% for Engineering and Technology. The average job rate per day exceeded for the first time 1.6 Million jobs/day, and the overall amount of CPU hours increased by 26.4% in the first year of the project.

The infrastructure provides both high throughput computing and cloud compute/storage capabilities by federating advanced computing services provided and funded at national level and by international research organizations, namely EMBL-EBI and CERN.

Two are the main technical platforms that support research virtual research environments: a federated cloud platform offering storage and compute IaaS services including to date 21 providers of which one is commercial, and a high-throughput data processing and analysis platform for data-intensive applications. Today EGI is the largest research advanced computing infrastructure for research worldwide in terms of geographical research and amount of aggregated storage and computing offered.

The EGI infrastructure is constantly evolving, both in terms of services provided and service management processes and tools for federated IT service management across the federation. Both the available resources (+23%) and the aggregate usage of the capacity (+20%) have considerably increased during 2015.

The overall quality of the service provided by EGI is good. Availability and reliability are in line with the previous years. The central services supporting the federation have been provided with no deviations, also considering the very high quality targets defined for these core tools.

Supported by the EGI-Engage project, EGI Operations have developed the operational policies and procedures to handle the new service types introduced with the federated cloud platform. The security policies have been aligned with the new types of services provided by EGI, and a security threats risk assessment has been contacted at the end of the project year. The assessment document is being updated at the moment of writing.

EGI is implementing the IT Service management processes based on the FitSM standard to all the production services identified by the newly revised service portfolio. As part of this programme, EGI has been establishing Service Level Agreements with the existing and prospective user communities through a new SLA negotiation procedure that helps defining the users service requirements and the committed service level requirements offered by EGI service providers. The new process was defined and introduced during PY1; today three service level agreements have been completed involving the following communities: the Swedish Bioinformatics Infrastructure for Life Sciences (ELIXIR Swedish head node, bioinformatics), the Distributed Research Infrastructure for Hydro-Meteorology lead by Fondazione CIMA (DRIHM) and the virtual research community lead by the Faculty of Science – Chemistry, Utrecht University (MoBRAIN, structural biology and medical imagining).

EGI has been taking a leading role in the international landscape of collaborating e-infrastructures, to support use cases from international research collaborations and to facilitate the coordinated service provisioning for large international research projects. EGI.eu signed a MoU with Compute Canada, which completes the list of collaboration agreements in place with e-Infrastructures in North and South America, Africa-Arabia, China, India and the Asia-Pacific region.

The EGI production infrastructure roadmap is moving along several paths. The first is strengthening the existing production services, enhancing the user experience in particular for the relatively new services of the federated cloud. Priorities for PY2 are the improvement of uniform use experience for the cloud users across the EGI federation, the definition of new OCCI configurations to lower the barriers in joining the federation by new cloud providers and the continual improvement of documentation. This will allow the cloud services of the federation to evolve with more reliable and consistent federation capabilities that can ensure high productivity for the users and less overhead for the service providers.

The capabilities of the services offered are being extended through the development efforts of EGI-Engage focused on existing and new technical platforms, and with the support of other collaborating projects like AARC (AAI) and Indigo-Datacloud (cloud data federation on cloud and brokering). The EGI Marketplace being prototyped in WP3 will allow the publishing and promotion of composite thematic services supported by EGI but operated by community PaaS providers. In controlling the publishing process, EGI will ensure the quality of the services of the Marketplace while integrating these services within its operations framework.

During PY15 20 different projects starting publishing cloud virtual appliances in the open community platform AppDB supporting the virtual appliance management capabilities as well as virtual image management capabilities. In 2016 Q1 the total number of virtual appliances shared in the library are 83.

1 Introduction

The EGI infrastructure builds on 15 years of design, development, and production deployment of geographically distributed data analysis services. The production infrastructure federates hundreds of resource centres, to serve thousands of users organised in hundreds of research communities. This document provides an overview of the production infrastructure, in terms of capacity and usage by the communities, and the activities that enable the federation.

Section 2 provides an overview of the status of the production infrastructure, the resources accessed by users, in terms of geographical distribution, deployed capacity and resource consumption, both for high throughput computing and cloud. The data and metrics provided by this section has been gathered from the EGI operational tools, such as the accounting and the monitoring services, and the trends in the last year, or several years, are analysed.

Section 3 focuses on the operational coordination and the operational framework that support the EGI federation, the processes, procedures and the central services that integrate the services operated by the resource centres and the operations centres to ensure the uniformity and quality of the service provisioning to EGI users. The section describes the evolution in the operational procedures, security coordination, service management and software quality assurance. The new services evaluated for production deployment need to be validated versus the operational processes and policies, and this process triggers – when needed- the evolution of the operational tools and the extension of the federation framework.

Section 4, describes the roadmap for the evolution of the EGI Operations in the coming months, with an overview of the new services, or the new access modes to existing services, that will have to be integrated in the production infrastructures.

The federation of cloud services, in production for more than one year, has continued. The capacity has expanded by integrating new sites, and the capabilities offered to the users have been extended, through the extension of standard interfaces and the integration of additional APIs including some native cloud management framework interfaces. The work done to improve the support for cloud services affected all the levels of the EGI operation, and it has been described in all the following sections.

2 The EGI Infrastructure

Being the EGI an international federation of data centres, multiple distributed resource and service providers contribute to the delivery of services. These can be either aimed at providing capabilities for the end-users, or to enable the federation of national infrastructures (NGIs) and ICT infrastructures operated by CERN and EMBL-EBI (European Research Infrastructure Organizations – EIROs). EGI federates individual capabilities while fully delegating the operations and policies for service access to the individual providers. By doing so, EGI leverages the specializations and competencies of EGI partners and opens up capabilities otherwise accessible just to local research communities. By doing so, international research collaborations and projects supported by multiple countries, find in EGI a natural environment where advanced computing capabilities are accessible to the entire collaboration, and in-house private infrastructures and services – where available - can be federated to join EGI at the same time and shared with other research communities where applicable.

Capital and operational expenditures for the delivery of EGI services are completely funded by the EGI participants. EGI-Engage contributes to fund the coordinated operations of the federation, so that a minimum set of interoperable service management processes and tools are maintained, innovated and operated. This ensures that national e-Infrastructures can be "interconnected" and transformed into a single international research system.

The high-performance analysis platform started its operations in 2004, while the federated cloud platform went into go-to-market stage in May 2014.

2.1 EGI Service Portfolio

The first edition of the EGI service portfolio was developed during 2013 to improve service orientation and clarify the unique offering that current and potential beneficiaries can request. This first version focused mainly on services internal to EGI as essential to enable the federation to work together and serve international research communities. This work was initiated in the context of improving the maturity in managing services by developing and implementing best practices for ensuring clarity of service offering and warranties and meeting the expectations of beneficiaries[1].

Following the improved maturity in designing and delivering services, the EGI service portfolio now covers both services that are internal to the EGI and services that EGI collectively delivers to the beneficiaries (researchers and SMEs/Industries).

In July 2015, EGI.eu has established the Services and Solutions Board (SSB) as a new body responsible for managing the portfolio of services and solutions regarding EGI.eu and the EGI

[1] http://fitsm.itemo.org/

[2] https://documents.egi.eu/document/2374

federated services, ensuring transparency across functions, and advising the EGI Council[2]. Following the creation of the SSB, the group has worked extensively to implement the service portfolio process (SPT) from FitSM, to define the templates and to update the EGI service portfolio. According to the established practice, each service is described in a Service Design and Transition Package (SDTP) document[3] composed of the following sections: value proposition, the business case, the service design, and the service transition plan. The expected impact of this activity includes:

- The improvement of service orientation
- The improvement of capabilities to promote EGI services and their value
- The improvement of management of services
- The alignment with the EGI strategy
- The management interoperability in federated environments
- A better understanding of all the components, dependencies and processes behind service delivery

The following table presents a summary view of the proposed update to the EGI service portfolio.

Table 1 The EGI service portfolio in 2016 Q1. The portfolio is continuously updated and managed under the control of the Services and Solutions Board and with the involvement of the EGI management. The innovation of the portfolio is lead by user innovation and by the opportunities created by the always-evolving technology market introducing new platform-driven solutions.

Service Category	Service name	EGI	EGI federation	Research
Compute ***Aimed at individual researchers, or national and international research collaborations that want to run their data- and computing-intensive experiments.***	Cloud Compute	√	√	√
	Cloud Container Compute	√	√	√
	High-Throughput Compute	√	√	√
Storage ***Targeted at researchers and research communities that need to access digital resources on a flexible environment.***	Object Storage	√	√	√
	File Storage	√	√	√
	Archive Storage	√	√	√
Data Management ***Aimed to help individual researchers, and research***	File Transfer	√	√	√
	Content Distribution	√	√	√

[2] https://documents.egi.eu/document/2374

[3] https://documents.egi.eu/document/2550

Service Category	Service name	EGI	EGI federation	Research
communities that have large-scale data management and computational capacity requirements.	Federated Data Manager	√	√	√
	Metadata Catalogue	√	√	√
Software and Service Platform ***Primarily aimed at Research Infrastructures and Resource Centers already within the EGI community or wishing to become part of it. It can also help other IT service providers that are geographically and/or structurally dispersed, and wish to organize themselves for federated service provision.***	Configuration Database	√	√	
	Accounting	√	√	
	Service Monitoring	√	√	
	Helpdesk	√	√	
	Attribute Management	√	√	
	Identity Provider Proxy	√	√	
	Marketplace	√	√	√
	Training Infrastructure	√	√	√
	Training Marketplace	√	√	
	Validated Software and Repository	√	√	
	Operations Tools	√	√	
	Virtual Research Environments	√	√	√
	Collaboration and Community Management Tools	√	√	
	Configuration Database	√	√	
Coordination and Support ***Primarily aimed at Research Infrastructures and Resource***	Project Management and Planning	√	√	

Service Category	Service name	EGI	EGI federation	Research
Centres already within the European Grid Infrastructure EGI community or wishing to become part of it.	Operations Coordination and Support	√	√	
	Technical Coordination	√	√	
	Security Coordination	√	√	
	Community Coordination and Development	√	√	
	Strategy and Policy Development	√		
	ITSM Coordination	√	√	
	Communications and Promotion	√	√	

The EGI services and solutions can be accessed through the following access policies:

- Policy-based: users are granted access based on policies defined by the EGI resource providers or by EGI.eu; such policies usually apply to resources being offered "free at point of use" to meet some national or EU level objective; for instance, a country may offer free at point of use resources to support national researchers involved in international collaborations.
- Wide access: users can freely access scientific data and digital services provided by EGI resource providers.
- Market-driven: users can negotiate a fee to access services either directly with EGI resource providers or indirectly with EGI.eu.

Services allowing access to rival resources (e.g. computing capacity or storage space) are usually provided under a policy-based or market-driven access policy. On the other hand, services allowing access to non-rival resources (e.g. software packages or scientific data) are usually provided under a wide access policy.

It may be that not all the access policies are available for each and every resource, service or scientific data set. Services and solutions are primarily intended for research purposes, but EGI is working on developing its business by tailoring existing services and developing new ones for the education sector and for the commercial exploitation.

From a service management point of view, EGI services are hosted by Resource Centres, the smallest resource administration domain in EGI. 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 the users whose access is granted by exposing common interfaces to the users.

2.2 EGI Operations

The resource providers manage and operate (directly or indirectly) all the operational services required to an agreed level of quality as required by the Resource Centres and their user community. They are also responsible for the maintenance, coordination and integration of the resource centres that build their individual e-infrastructures. The Resource infrastructure Providers liaise locally with the Resource Centre Operations Managers, and represent the Resource Centres within EGI.

The EGI resource infrastructure providers are:

- **Council-members resource providers:** National Grid Initiatives (NGIs) and European Intergovernmental Research Organizations (EIROs), who are represented in the EGI Council and directly contribute to the sustainability of EGI.
- **Integrated resource providers:** International organizations who contribute resources and have a collaboration agreement with EGI through a Memorandum of Understanding:
  - Asia Pacific Region (including resources from Australia, China, India, Iran, Japan, Malaysia, New Zealand, Pakistan, Philippines, South Korea, Taiwan, Thailand, Vietnam)
  - "Africa Arabia" (including resources from South Africa, Egypt, Morocco, Tanzania, Kenya, Nigeria, Algeria, Senegal, Ethiopia, Tunisia, Ghana)
  - Ukrainian National Grid (UNG)
  - Latin America (including resources from Brazil, Chile, Mexico)
  - IHEP (China)
  - Canada
- **Collaborating resource providers:** Other international organization who have strong collaborations with EGI:
  - Open Science Grid (USA)
  - ComputeCanada (Canada)
  - C-DAC (India)

The following map shows the international e-Infrastructure and the representative partners EGI is collaborating with.

2.3 Status

In February 2016 EGI comprises resources provided across **58 countries and 2 European Intergovernmental Research Institute (CERN and EMBL)**, of which **26 are EGI Council members**. From an operational point of view there is no difference between the integrated resources and the Resource Centres from the EGI council participants, all these receive the same level of support (this ensures the possibility to deliver professional services to third countries and developing ones). EGI Resource Centres, and must fulfil a minimum set of requirements and accept EGI policies and service management procedures.

The peer resource providers are infrastructures with which EGI has interoperations agreements, and common user communities, but do not share the EGI operational infrastructure.

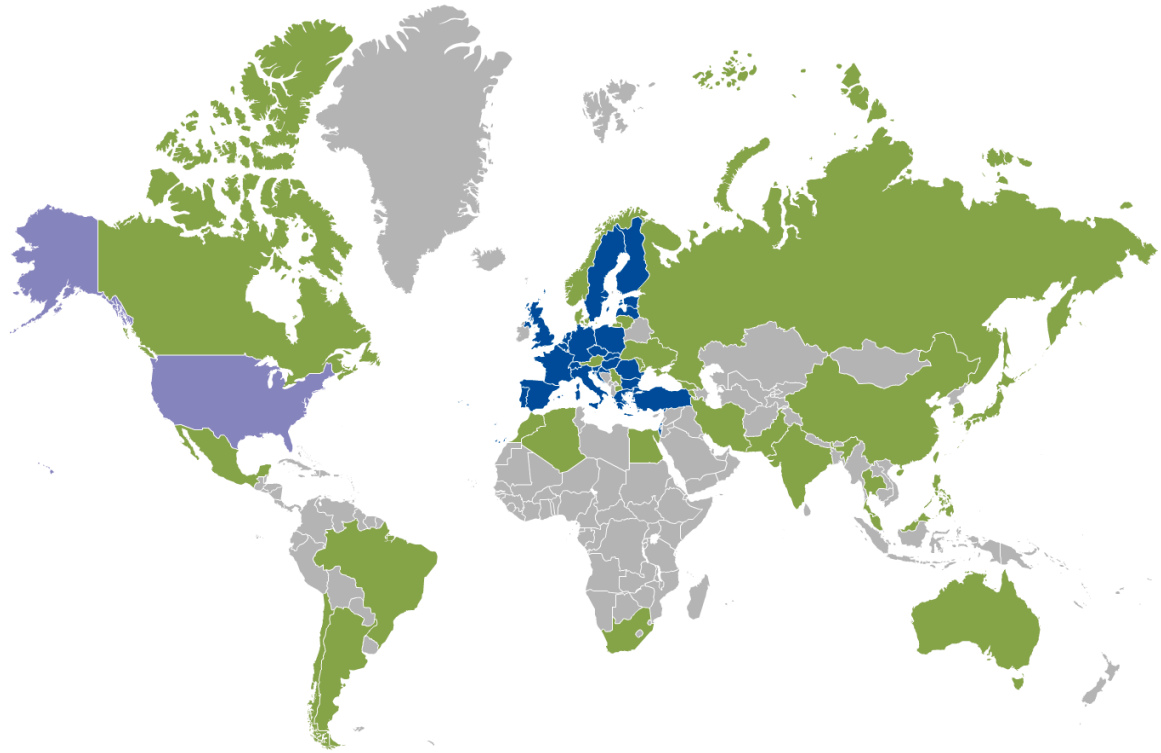

Figure 1 Worldwide presence of EGI Resource Centres: EGI council members in blue, integrated RPs in green, collaborating RP in purple.

EGI operational procedures are being revised with the purpose of simplifying them and facilitating the integration of new resource providers while reducing the human effort for operating the national infrastructures.

2.4 Distribution of capacity

The EGI national participants – the NGIs – are organisations set up to manage the resources provided in their countries by the resource centres to the EGI. They represent the country's single

point of contact for EGI as well as to liaise with government, research communities and resource centres as regards ICT services for e-Science.

Each NGI operations are supported by an Operations Centre, defined as a centre offering operations services on behalf of the Resource infrastructure Provider, and it can serve multiple RPs. Examples of these services are supporting the sites in the certification process, deploying the monitor services at NGI level or information system, and liaise with EGI during the software upgrade campaigns.

EGI currently comprises 27 national operations centres and 7 federated operations centres encompassing multiple NGIs. The federated centres in Europe NGI_IBERGRID, NGI_NL and NGI_IT, each containing two countries, are the result of a collaboration agreement that is expected to continue in the next PYs. In contrast, integrated federated centres in Asia Pacific and Latin America encompass a large number of countries, as in those regions Resource Centres are sparse and their number does not justify the overhead for the creation of a national operations centre, but suggests that an international collaboration is in place. The creation of new NGIs in those regions will depend on their expansion plans and on national policies.

The following table shows the distribution of resource centres per country and per operation centre.

Table 2 shows the resource centres number per country and per operation centre.

Operations Centre	Country (RCs Number)	Resource Centres
AfricaArabia	Algeria (1), Egypt (1), Morocco (2), South Africa (5)	9
AsiaPacific	Australia (1), China (1), India (2), Iran (1), Japan (2), Malaysia (3), Pakistan (2), South Korea (4), Taiwan (6), Thailand (4)	26
CERN	Switzerland	1
IDGF[4]	Hungary	1
NGI_AEGIS	Serbia	6
NGI_ARMGRID	Armenia	1
NGI_BG	Bulgaria	2
NGI_CH	Switzerland	5
NGI_CHINA	China	1
NGI_CZ	Czech Republic	3
NGI_DE	Germany	18
NGI_FI	Finland	10

[4] At the moment of writing the status of the desktop grid operations centre – offering volunteer computing – is on hold, since the leading institution asked for decommissioning but the procedure has not been finalized yet.

Operations Centre	Country (RCs Number)	Resource Centres
NGI_FRANCE	France	17
NGI_GE	Georgia	1
NGI_GRNET	Greece	12
NGI_HR	Croatia	4
NGI_HU	Hungary	2
NGI_IBERGRID	Portugal (5), Spain (16)	21
NGI_IL	Israel	5
NGI_IT	Austria (2), Italy (45)	47
NGI_MARGI	FYROM	2
NGI_MD	Moldova	2
NGI_NDGF	Denmark (1), Estonia (2), Finland (1), Latvia (2), Lithuania (1), Norway (1), Sweden (2)	10
NGI_NL	Belgium (3), Netherlands (14)	17
NGI_PL	Poland	13
NGI_RO	Romania	10
NGI_SI	Slovenia	2
NGI_SK	Slovakia	8
NGI_TR	Turkey	3
NGI_UA	Ukraine	15
NGI_UK	United Kingdom	24
ROC_Canada	Canada	8
ROC_LA	Brazil (3), Chile (3), Mexico (3)	9
Russia	Russia	9
Total OCs: 34	**Total Countries: 53**	**Total RCs: 324**

As shown above the total number of certified RCs in February 2016 amounts to 324.

In December 2014 there were 352 certified RCs: this decrease may be explained by the fact that the especially for the RCs with a low amounts of resources, both in terms of hardware and personnel, it was difficult maintaining a level of service in accordance to the EGI operational level agreement[5], so they have been suspended when the issues they are facing require more time to be solved, than it was allowed. Other RCs instead were decommissioned because they ceased operations. Currently 75 RCs are in the status of suspended or uncertified on GOCDB, and some of them are already in the process of being re-certified. The suspension of a resource centre is needed to ensure that the services that users expect to be production ensure the level of quality required to be used productively.

[5] https://documents.egi.eu/document/31

We usually distinguish between two categories of services: HTC and Cloud. There are also some RCs that provide both categories of service, so they are counted in both of them (Table 1).

In the following sections they will be analysed separately.

Table 1 Number of HTC Resource Centres and Cloud providers (February 2016)

Number of HTC resource centres	303
Number of cloud providers	21
Number of mixed HTC/cloud providers (also included in the totals above)	7

2.5 High Throughput Computing

The HTC services offered by a resource centre can be grouped in two categories:

- Grid compute: allows users to run computational tasks on high quality IT resources, accessible via a standard interface and supporting authentication/authorization based on a membership within a virtual organization.
- Grid storage: allows files to be stored in and retrieved from high quality IT resources, accessible via a standard interface and supporting authentication/authorization based on a membership within a virtual organization.

According to the OLA[6], each site may provide one or both of this kind of services. The HTC compute platforms supported in EGI are: ARC-CE, CREAM, UNICORE and GLOBUS. Since the UNICORE and GLOBUS resources are not published in the information system, for them we can provide only the number of the certified instances registered in the GOC-DB.
The total grid compute capability is shown in the Table 2.

Table 2 EGI grid compute capacity (February 2016)

	Logical cores	HEP-SPEC 06
2014 (December)	527248	4211709,28
2016 (February)	651748	5841854,65
Yearly increase	23.61%	38,71%

At the end of 2014 the total amount of logical cores was 527248, so the relative increase exceeds 23%. This is in-line with the trends of the last years, as shown in the following graph.

[6] https://documents.egi.eu/document/31

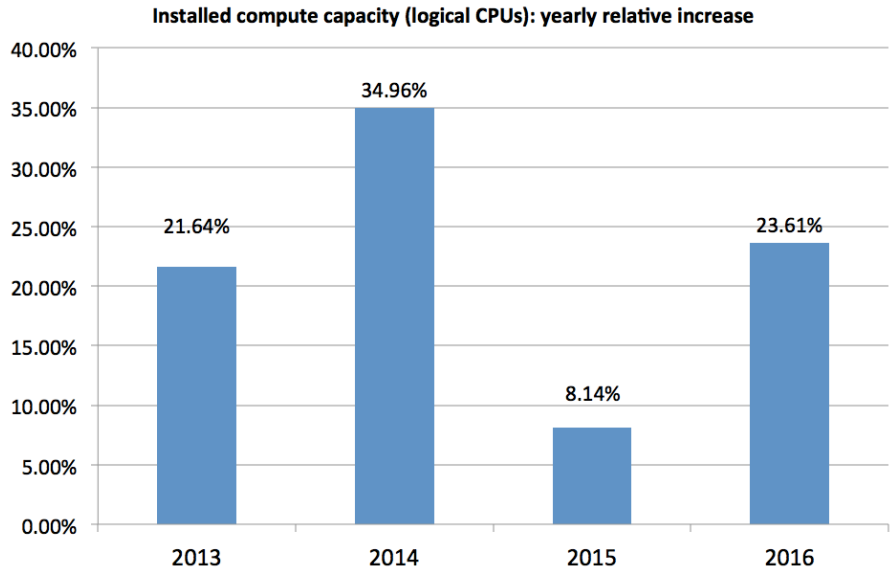

Figure 1 Yearly relative increase of the number of installed logical CPUs in the EGI HTC platform (2013-2016). As shown in this diagram, the yearly increase is very variable; however, in 2016 the relative increase shows a significant increase when compared to 2015.

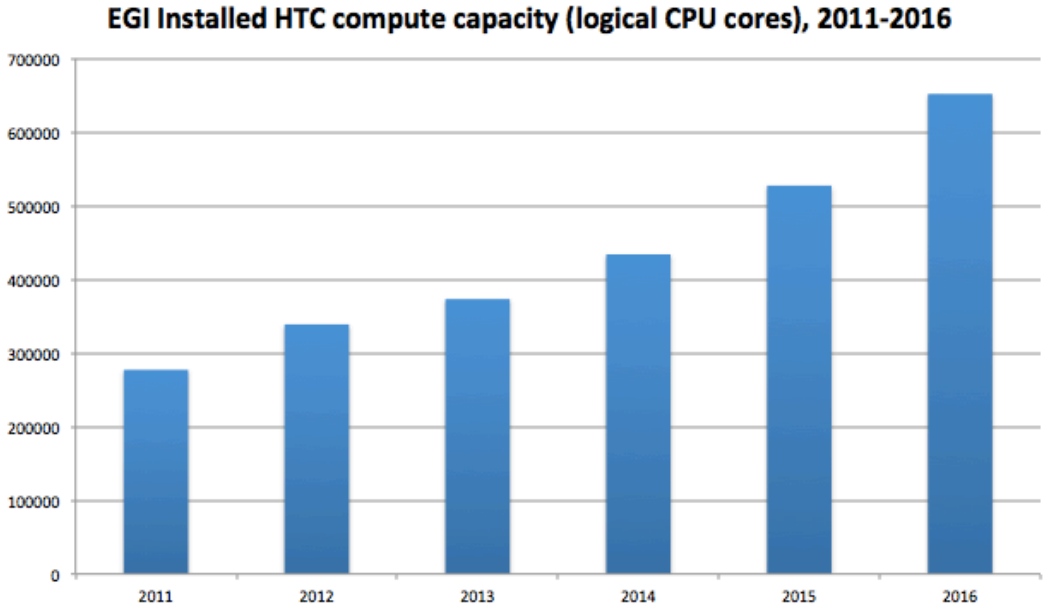

Figure 2 In 2016 the installed capacity of EGI (High Throughput Computing platform) broke the wall of 650,000 installed logical CPU cores.

The total number of RCs slightly decreased but the size of the certified ones grew in terms of resources contributed, thus increasing the total capacity of the infrastructure. As shown in Figure 3, where the logical cores and HEP-spec06[7] power distribution is shown in as well.

[7] https://w3.hepix.org/benchmarks/doku.php

The increase in the capacity of the EGI infrastructure is driven by a number of factors, which include for example the national infrastructure development plans and the needs of the communities who federate === resources federated in the EGI.

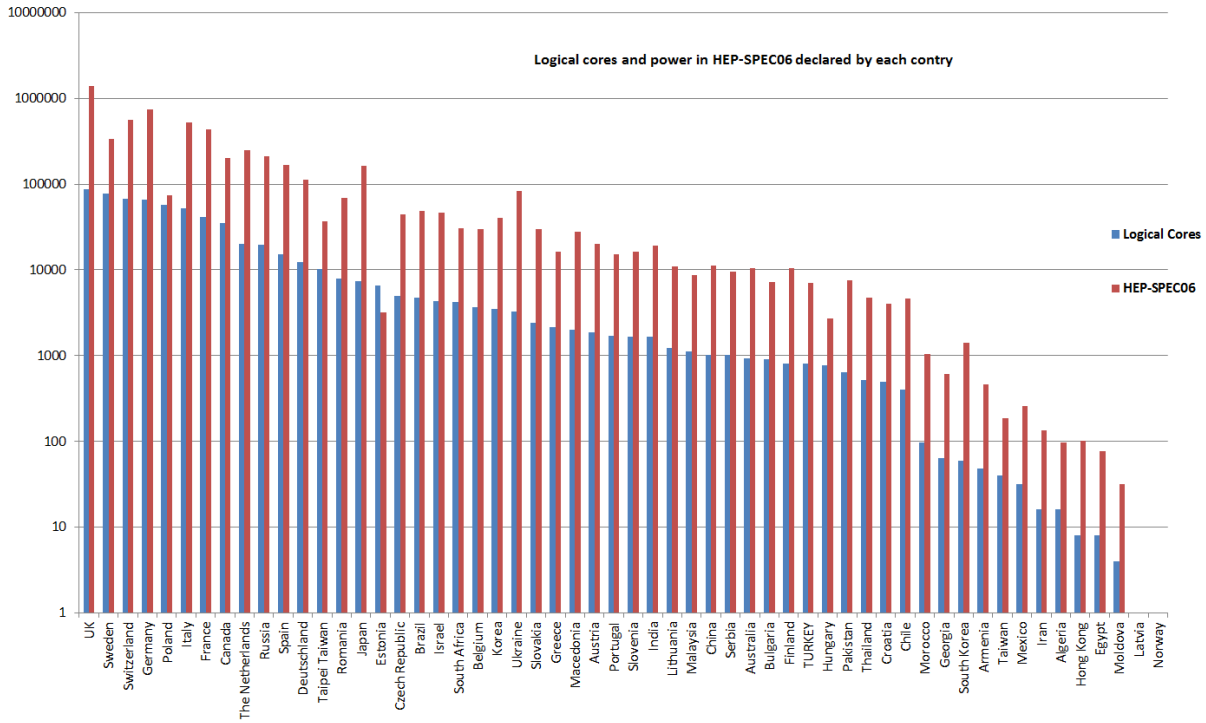

Figure 3 Logical cores and hep-spec power provided by each country.

Figure 4 shows the distribution of computing element types across the NGIs. HTC Computing is provided through at least 5 different type of middleware, and this diversity allows EGI to support different use case by exploiting the different features provided by the software available.

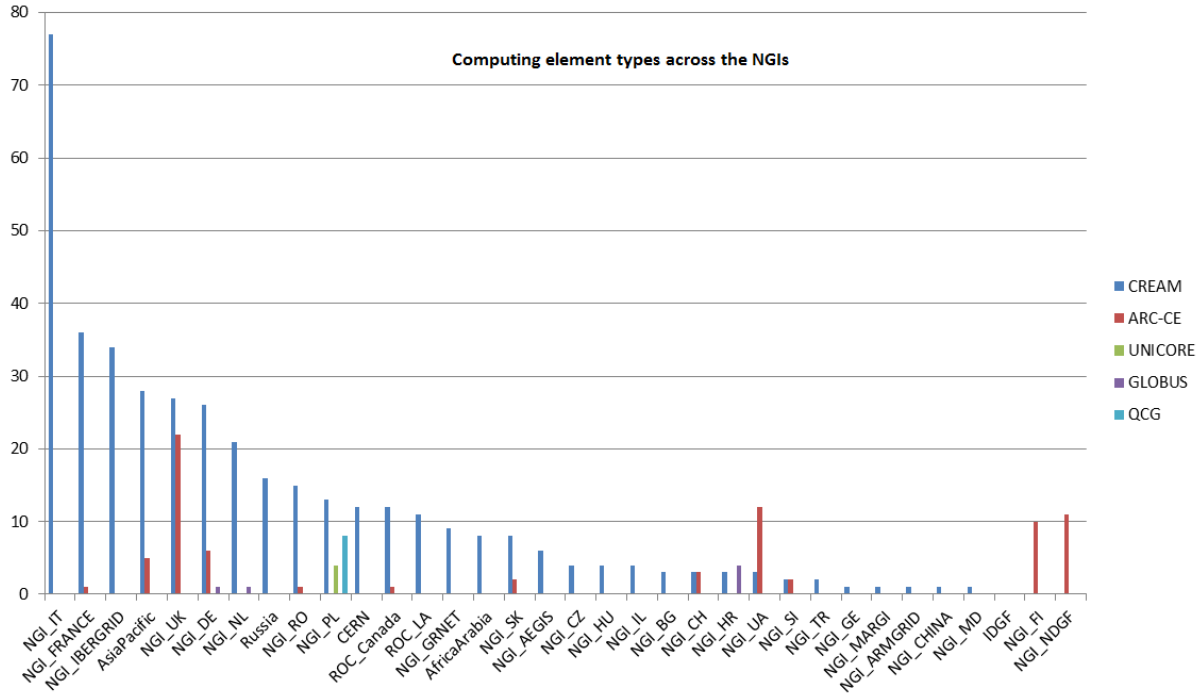

Figure 4 Computing Elements distribution across the NGIs grouped by type of interface (CREAM, ARC-CE, UNICORE Globus and QCG)

The EGI infrastructure provides also compute resources for parallel jobs. The numbers of resource centres that support parallel computing via MPI jobs are 54 as results in February 2016 (Figure 5), or rather 67 computing elements in total. In 2014 there were 76 RCs supporting MPI: as explained above, MPI computing was concentrated on a fewer larger Resource Centres, and the overall capacity offered considerably increased.

Information about MPI capabilities is not only published by services via the Information Discovery Service, but they are also registered into the EGI service registration facility GOCDB. In addition, during 2015 a new accounting publisher was deployed, this new release of accounting is capable of reporting accounting information of multi-core jobs, where computation is parallelized by running concurrent threads on different cores. As accounting records are accumulated over time, MPI accounting capability will be a more accurate indicator of the amount of parallel computing workload supported by EGI, and will also complement the information about MPI support available in GOCDB and the information system; this type of new accounting has been adopted by an increasing number of Resource Centres during 2015.

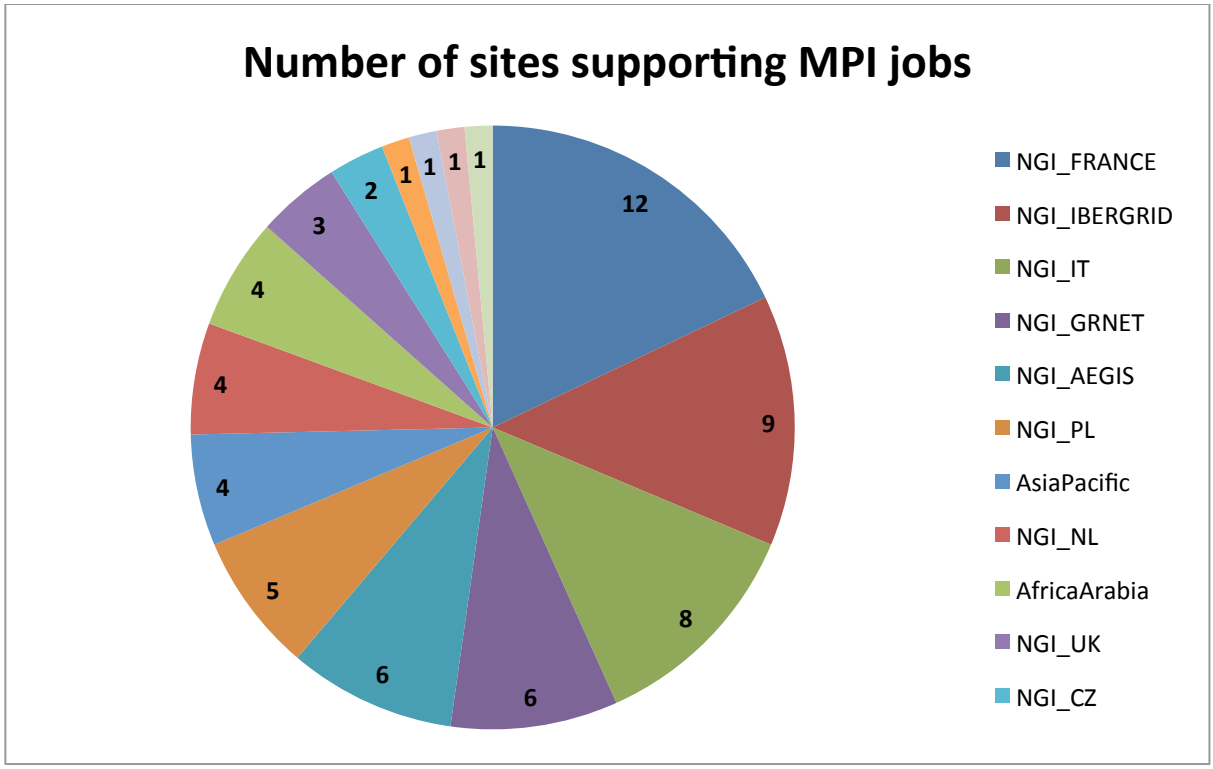

Figure 5 Number of EGI RCs supporting MPI jobs (February 2016, source: GOC-DB)

Three main data management products are available in EGI to provide access to geographically distributed data: DPM, dCache and STORM.

The total amount of storage certified service end-points is 313, which corresponds to a total disk capacity of about 264.18 PB. In December 2014 the total disk capacity reported[8] was 236.19 PB, so it increased by 11.85%. Instead the total tape capacity (also called nearline storage), which is mainly provided by CERN and WLCG Tier-1 RCs amounts to 239.8 PB. In April 2014 the corresponding value was 168.8 PB, so the increase was 42.06%.

The distribution of disk storage resources among the EGI operations centres is shown in the following figure (Figure 6), which shows that the disk capacity is concentrated across five NGIs: NGI_IT, NGI_UK, NGI_DE, NGI_IBERGRID and the Asia-Pacific region in descending order.

[8] In the GSTAT monitoring tool http://gstat.egi.eu/

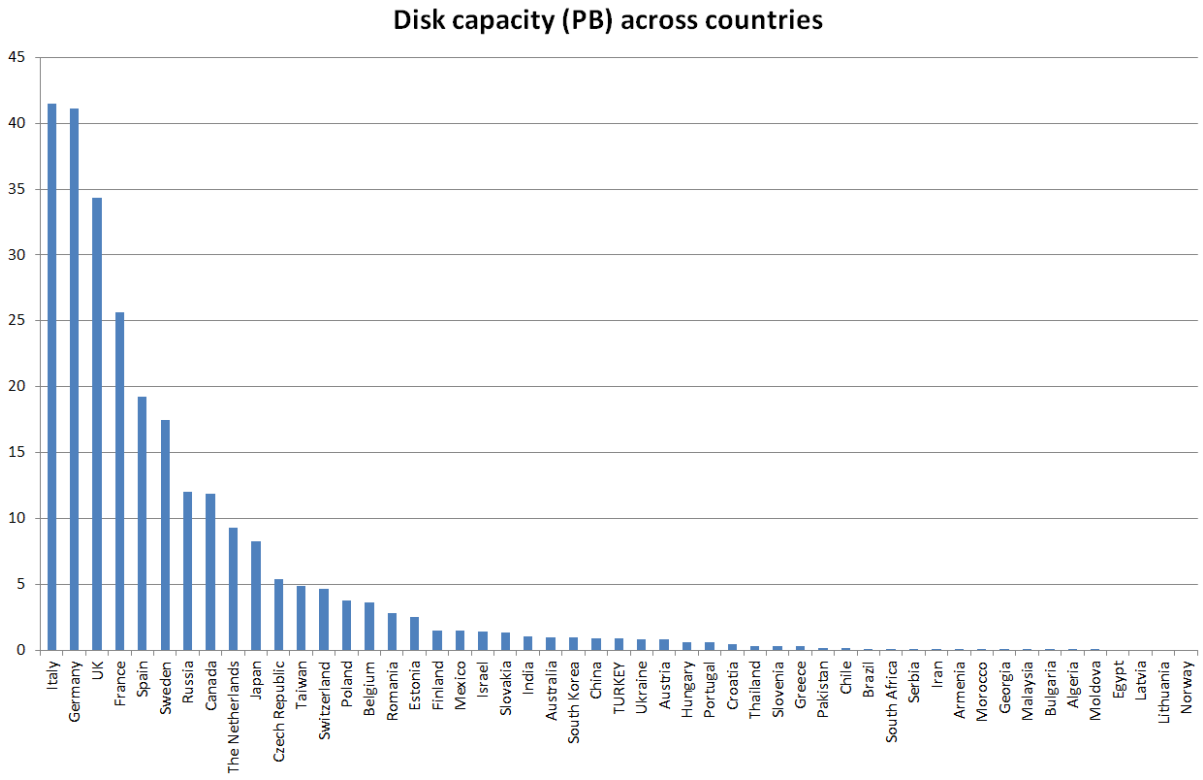

Figure 6 Disk capacity distribution across the NGIs (source: GSTAT)

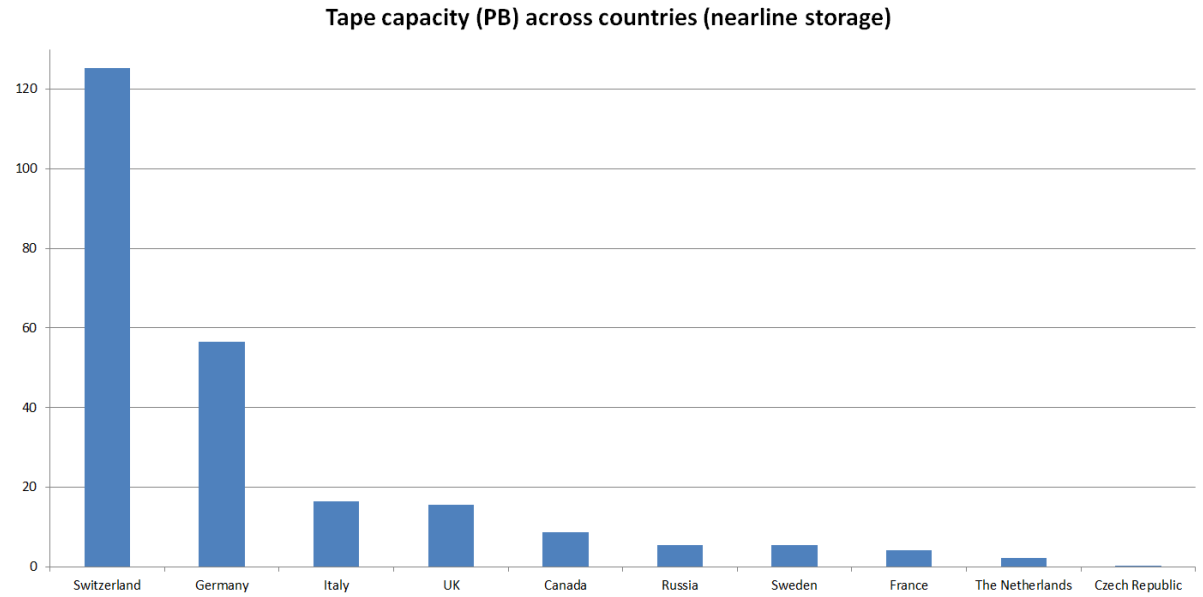

Figure 7 Tape capacity (online storage) distribution in PB across the NGIs and CERN (Switzerland region) (source: GSTAT)

2.6 EGI Federated Cloud

The EGI Federated Cloud is part of EGI production infrastructure, a seamless HTC of academic private clouds and virtualised resources, built around open standards and focusing on the requirements of the scientific community. It is **in production** since mid-May 2014.

The Federated Cloud is targeted at researchers and research communities that need to access digital resources on a flexible environment, using common standards to support their data- and computing intensive experiments.

2.6.1 Cloud federation

The impact of cloud computing on science, scientific development and education is undoubtedly increasing. Researchers and research institutes, projects, communities turn to clouds more and more often when they need a platform to store, share, process or archive large research data in a reliable and user-friendly way. The number of cloud services and service providers who target specifically the research and educational sectors is growing. OpenStack is becoming the de-facto standard for building both private and public clouds for this domain. Current trends indicate that in the future research communities will demand capabilities to federate services from multiple providers in order to support complex, community-specific, cross-institutional or international use cases.

Given its long-running experience in federating IT services for research and education, EGI is ideally positioned to be a key player in the federated cloud computing landscape and specialize on building federated clouds for research and educational use cases.

The EGI Federated Cloud collaboration includes technology providers, resource/cloud providers, and user support and system administration personnel from various communities, including the EGI community itself. The collaboration members:

- Identify and integrate open source tools and services that enable cloud federations for research and education.
- Develop and maintain of tools and services to fill gaps in third party solutions to reach production quality cloud federations.
- Provide consultation and training for communities on how to build a federated cloud to meet custom community demands under certain constraints.
- Provide training and support for existing and potential users of cloud federations about topics, such as how to port or develop cloud-based applications; how to operate services in the cloud, how to join a cloud federation with a service.
- Facilitate the reuse of cloud federation tools and services across participating cloud federations to lower total cost of development and to improve cloud sustainability.

- Promote Platform as a Service (PaaS) and Software as a Service (SaaS) environments that are proven to be robust and reusable across communities to interact with federated IaaS clouds.
- Provide service management and security oversight for participating clouds and cloud federations.
- Act as a discussion forum where cloud federations can be discussed and specific questions can be analysed with top-world experts.
- Organise dissemination and marketing events, workshops and conferences relating to the topics of the collaboration.

Joining the EGI Federated Cloud collaboration can bring various benefits:

- Participate in the review and selection of tools and technologies that enable cloud federations for research and education.
- Participate in software development and integration projects that build solutions for cloud federations.
- Become or act as an advisor towards scientific communities about building federated clouds.
- Became or act as a trainer and consultant of users who want to use federated clouds offered by members of the collaboration.
- Offer technological solutions (IaaS, PaaS, SaaS) that can be used in various federated clouds to satisfy demands in the research and education sectors.
- Become a service management and/or security expert specialised on clouds and cloud federations and contribute to the work of the security and operational oversight group.
- Promote community-specific, but reusable federation solutions and approaches to other communities within and beyond the EGI collaboration.

2.6.1.1 Federating clouds

A cloud federation is an interconnected cloud environment between two or more service providers. Such setups are typically motivated by:

- **Distributed capacities or capabilities**: A single cloud provider cannot provide all the capacity or all types of cloud services that the community requires. Cloud bursting is a typical example for such a setup.
- **Restrictive data policies**: Certain policies may restrict for an institute or for a community as a whole to move data from its current location. Data providers must therefore have to become cloud providers to allow users to send processing algorithms to the data - in the form of virtual machine images.

- **'Too large' data sets**: Certain datasets may be too large to move from their existing locations. Data providers must therefore become cloud providers and allow users to send processing algorithms to the data.
- **Distributed investments**: A community may decide to build cloud services (or purchase cloud services) at multiple sites - for example to stimulate local economies.
- **Distributed expertise**: A community may decide to procure from multiple cloud providers to create knowledge hubs at multiple locations.

Scientific communities can have different motivations for building/using federated clouds. They can be at different stage in implementing a federated cloud; moreover they may be already committed to certain technological solutions that need to be incorporated into their cloud federation. One federation approach cannot fit all; therefore the EGI Federated Cloud collaboration needs to provide ad-hoc solutions for building cloud federations.

2.6.1.2 Services in a federated cloud

What makes a cloud federation? Cloud services - IaaS, PaaS or SaaS - are a necessity. But besides this, there is also the need for services that interconnect these cloud environments. Despite the large diversity in the type of cloud federations, a relatively small number of building blocks (or federator services) can be identified in almost all of them. Federation models can be loose of tightly coupled depending on the number of common requirements the members of the federation need to comply to. The table below shows how different cloud infrastructures can be federated through EGI services.

Table 3 EGI services for the implementation of a cloud federation.

Federator service	Capabilities offered to the federation	EGI product
Service registry	Having a registry where all the federated cloud sites and services are registered and state their capabilities. The registry often provides the 'big picture view' about the federation for both human users and online services (such as service monitors).	GOCDB
Information system	Having a database (often with a web interface) that provides real-time view about the actual capabilities and load of federation participants. The information system can be used by both human users and online services.	BDII
Virtual Machine image catalogue	Having a catalogue of Virtual Machine images (VMIs) that are usable by the IaaS cloud providers and encapsulate	AppDB

Federator service	Capabilities offered to the federation	EGI product
	those software configurations that are useful and relevant for the given community (typically, pre-configured scientific models and algorithms). To maximise usability of VMIs across cloud sites the images should be in a format that's supported at every federation member site (or at least that can be converted to such formats)	
Image replication mechanism	Having a system that automatically replicates VMIs from the VMI catalogue to the federation member sites, as well as removes them when needed. Automated replication can ensure consistency of capabilities across providers and is very often coupled with a VMI vetting process to ensure that only properly working, and relevant VMIs are replicated to the cloud sites of the community.	VMCaster, VMCatcher
Single sign-on for users	Allow users of the federated services to register for access only once while accessing distributed trusted services. Single sign-on is increasingly implemented in the form of identity federations in both industry and academia.	X509 Certificates from IDGF and PERUN
Integrated view about resource/service usage	Complete usage (accounting) reports from the federated providers, integrating information and presenting it in such a way that both individual users as well as whole communities can monitor their own resource/service usage across the whole federation.	APEL accounting system and portal
Integrated interfaces or user environments	Having interfaces through which users and user applications can interact with the services offered by the various cloud providers. In case of an IaaS cloud federation these interfaces offer compute, storage and network management capabilities. The interfaces can be harmonised across all participating cloud providers - in which case the providers are responsible for implementing the agreed standard - or can be native at	OCCI API and rOCCI client

Federator service	Capabilities offered to the federation	EGI product
	the different sites. In this latter case centrally maintained user environment or portals can hide heterogeneity from the users and can translate user requests to diverse native formats.	
Integrated helpdesk and user support	Having a helpdesk and user support process that offers a single point of entry for users and ensures issue resolution on a timely manner despite the distributed, federated landscape.	GGUS with national and topical support teams
Offloading excess workloads		
Shared operational practices	The participating service/resource providers may share certain operational tools and practices at the level of the federation, for example use a shared system to collect availability and reliability statistics about their site, or to share and respond to security alerts.	ARGO monitoring system; EGI Operation teams at national and site levels

2.6.2 Cloud model

The EGI Cloud federation is a hybrid cloud composed by public, community and private clouds, all supported by the EGI Core Infrastructure Platform (AAI, Service Registry, Accounting, Monitoring and Federated Service Management). The EGI Federated Cloud is composed by multiple "realms", each realm having homogeneous cloud management interfaces and capabilities. A Community Platform provides community-specific data, tools and applications and can be supported by one or more realms.

EGI Cloud Federation: hybrid cloud (private, community, public).

EGI Cloud Realm: subset of cloud providers exposing homogeneous cloud management interfaces and capabilities. The Open Standards Cloud Realm supports OCCI and CDMI.

Community Platform: set of community-specific data, tools, applications, and brokering tools, which can be supported by one or more realms of the federation.

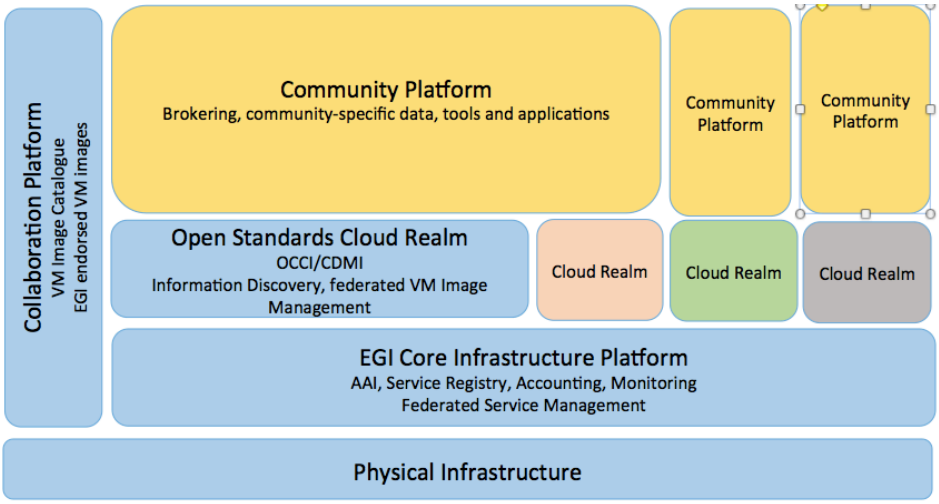

Figure 8 Platform architecture (EGI Core Infrastructure Platform, Cloud Realms and Community Platforms)

The EGI Cloud Federation offers different types of realm federation model, the requirements on the cloud providers vary depending on the type of model of choice (see Table 4).

Table 4 Requirements for cloud providers in order to be part of the EGI Federated Cloud. Various federation models apply depending on the model of choice.

EGI Cloud Federation requirements for cloud providers			
Requirements	**EGI Cloud Realms**		**Peer Realms**
	Open Standards Realm	**Other Realms**	
EGI AAI compliance	yes	yes	yes
EGI federated service management (processes, activities and policies) adoption	yes	yes	yes
Service registry	yes	yes	optional
Accounting	yes	yes	optional
Monitoring	yes	yes	optional
Information discovery	yes	optional	optional
IaaS open standards compliance - OCCI, CDMI	yes	optional	optional
VM image catalogue	optional	optional	optional
federated VM image management	yes	optional	optional
CMW native interfaces	optional	optional	optional
EGI endorsed VM images	optional	optional	optional

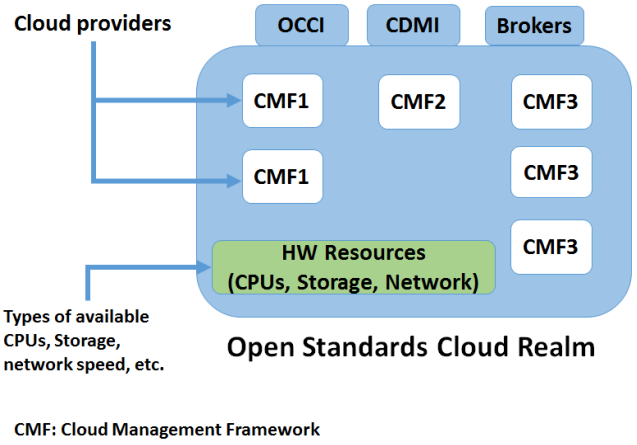

Capabilities	
Services	**Resources**
EGI AAI	Compute
Fed. Services Management	Block Storage
Service Registry	Object Storage
Accounting	Brokers and Orchestrators
Monitoring	High Speed Storage
Information Disc.	
IaaS OCCI & CDMI interface	
VM Image Catalogue	
Fed. Image Management	
EGI Endorsed VM image	

Figure 9 Requirements and configuration model of a cloud provider contributing the Open Standards cloud realm. In order to be part of this federation, the OCCI and CDMI open standard interfaces need to be exposed to ensure full portability.

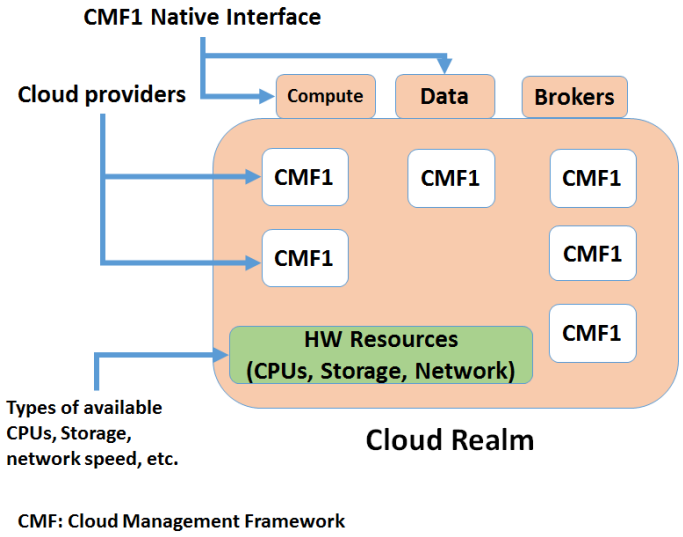

Capabilities	
Services	**Resources**
EGI AAI	Compute
Fed. Services Management	Block Storage
Service Registry	Object Storage
Accounting	Brokers and Orchestrators
Monitoring	GPGPU
Optional Service 1	High Speed Network
Optional Service 2	...
Optional Service N	

Figure 10 Requirements for a cloud provider joining a cloud realm of choice. In this case users interact via CMF native interfaces, so that portability is only ensured within the realm.

2.6.3 Cloud Infrastructure

Several resource centres joined to the Federated Cloud since its beginning: at the end of PY1 (beginning of February 2016) 21 RCs are part of the cloud infrastructure (Table 5) and other providers, once undergone to the certification procedure[9], are expected to join in (Table 6). From Table 5 we can see how the Federated cloud activity is led, in terms of resource provisioning, by IBERGRID, NGI_IT, NGI_SK and NGI_DE.

EGI cloud provider offers at least one the following service types:

[9] https://wiki.egi.eu/wiki/PROC09

- Cloud computes IaaS: including Virtual Machine (VM) management, block storage, a Virtual Appliance catalogue, and software and data distribution). The service allows scientists to manage VMs on demand, with customizable set of hardware, network and storage resources.

Cloud storage: block storage10 allows end-users and service providers to store files, images and other generic objects that can be accessed from any device with integrated basic processing capabilities.

The unique value proposition of the EGI federated cloud is the portability of applications across different providers, which make the environment particularly suitable for data-driven modern research applications based on data virtualization. The federated cloud is suitable for the processing and analysis of large data sets by bringing VM images to where the data resides (when online access to data or moving data do not scale) and for accessing data whose hosting cannot be accommodated by other countries/organizations for legal reasons. Both use cases are being addressed by the technical developments of EGI-Engage WP4 – the open data platform, which aims at implementing a federated data infrastructure suitable for publishing, using and reusing open distributed datasets.

Resource centres are free to use any Cloud Management Framework with the requirement that the CMF exposes interfaces compliant to the FedCloud standards[11].

Common CMF used are OpenStack and OpenNebula, but also Synnefo is supported in the federation. The common interfaces provided to access the virtualized resources are Open Cloud Computing Interface (OCCI) and Cloud Data Management Interface (CDMI).

Table 5 EGI cloud providers (2016 Q1)

Resource Centre	NGI	Number of cores declared	Amount of disk space declared	Cloud Management Framework
100IT	NGI UK	120	16 TB	OpenStack
BIFI	NGI IBERGRID	720	36 TB	OpenStack
CESGA	NGI IBERGRID	448	6 TB	OpenStack
CESNET-MetaCloud	NGI CZ	416	56 TB	OpenNebula
CETA-GRID	NGI IBERGRID	184	5 TB	OpenStack
CYFRONET-CLOUD	NGI PL	200	20 TB	OpenStack
FZJ	NGI DE	216	50 TB	OpenStack

[10] Solutions for federated object storage management as additional new capabilities are being tested.

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

GoeGrid	NGI DE	192	40 TB	OpenNebula
HG-09-Okeanos-Cloud	NGI GRNET	70	1 TB	Synnefo
IFCA-LCG2	NGI IBERGRID	2288		OpenNebula
IISAS-FedCloud	NGI SK	176	50 TB	OpenStack
IISAS-GPUCloud	NGI SK	96	6 TB	OpenStack
IN2P3-IRES	NGI FRANCE	192	5 TB	OpenStack
INFN-CATANIA-NEBULA	NGI IT	16	5 TB	OpenNebula
INFN-CATANIA-STACK	NGI IT	16	16 TB	OpenStack
INFN-PADOVA-STACK	NGI IT	144	5 TB	OpenStack
MK-04-FINKICLOUD	NGI MK	100	1 TB	OpenNebula
NCG-INGRID-PT	NGI IBERGRID	80	3 TB	OpenStack
PRISMA-INFN-BARI	NGI IT	300	50 TB	OpenStack
TR-FC1-ULAKBIM	NGI TR	336	40 TB	OpenStack
UPV-GRyCAP	NGI IBERGRID	128	5 TB	OpenNebula

Table 6 Cloud providers in the certification phase

Cloud providers under integration
Italy/RECAS-BARI**: new name of the cloud provider "PRISMA-INFN-BARI" just appointed (resources are being contributed and shared leveraging the investments of regional structural funds for the south of Italy)**
Germany/SCAI**: moving its resources from HTC to CLOUD**
Italy/INDIGO-CATANIA-STACK**: new Resource Centre**
UK/EMBL-EBI**: integration of a OpenStack infrastructure with OCCI is in progress**
Sweden/SNIC**: exchange of information on requirements for integrating the national cloud infrastructure**

In a distributed, federated Cloud infrastructure, users will often face the situation of efficiently managing and distributing their VM Images across multiple resource providers. Users need a catalogue of Virtual Machine images (VMIs) that are usable on the IaaS cloud provider sites and encapsulate those software configurations that are useful and relevant for the given community. (typically pre-configured scientific models and algorithms). To maximise usability of VMIs across cloud sites the images should be in a format that's supported at every federation member site (or at least can be converted to such formats). Users also need a system that automatically replicates VMIs from the VMI catalogue to the federation member sites, keeps them updated or removes them when not needed anymore. Automated replication can ensure consistency of capabilities across sites and it is being coupled with a VMI vetting process to ensure that only properly working and relevant VMIs are replicated to the cloud sites of the community.

The EGI AppDB service[12] has been extended to provide a Virtual Appliance Marketplace that stores metadata and access information on virtual machine images designed to run on a given virtualization platform.

AppDB's Virtual Appliance Marketplace provides the ground for managing and publishing versioned repositories of virtual appliances, in a way that integrates with the existing HEPiX VMCaster and VMCatcher[13] framework, currently in use by the EGI. Research Communities are responsible for creating and updating VM Images stored the Research Community, and for publishing a VM Image list using AppDB. Federated Cloud Providers use these lists to make the VMI available for instantiation at site level for the supported VOs, this process is automated through VMCaster/VMCatcher.

In a federated environment, brokering of computation to data requires information about the cloud resources being available in the federated environment. In the EGI federated cloud we are information published in the BDII (a LDAP-based service pulling information from service end-points). We are adopting the GLUE2 standard[14], and actively evolving it the GLUE2 working group of the Open Grid Forum to further extend the schema (v 2.1) to represent Cloud Computing, Storage and in the future Platform and Software services. The proposed extensions are currently under discussion in the working group.

All cloud services are registered in GOCDB (the EGI service registry) and are monitored with a cloud-specific instance of the EGI monitoring infrastructure based on Nagios: the monitoring service is an instance of the Service Availability Monitoring (SAM) production distribution, which features a set of cloud-specific probes. The replacement of SAM with the second generation product version of SAM named “ARGO” is planned.

2.7 Capacity consumption

EGI accounting information is gathered and stored centrally and accessible through the accounting portal[15]. Accounting information is aggregated by Operations Centre, whose list is obtained from GOCDB.

[12] https://appdb.egi.eu/

[13] https://github.com/hepix-virtualisation/vmcaster

[14] http://glue20.web.cern.ch/glue20/

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

Table 7 HTC compute resource usage in the last three years

	2015	2014	2013
Total normalized CPU time consumed (Billion HEP-SPEC 06 hours)	20.56	16.27	14.62
Total number of jobs (Million)	584.9	535	522.8
Average number of jobs per day (Million)	1.60	1.47	1.43

The overall quantity of HTC computing resources used in 2015 amounts to 20.56 Billion HEP-SPEC 06 Hours as shown in

Table 7, with an increment of 26% from 2014 (the increment in the 2014 compared to 2013 was 11%). The total number of jobs executed on the infrastructure is 584.9 Million, which corresponds to an average 1.60 Million job/day.

The increase of the normalized CPU time registered in 2015 (20%) higher than the increased number of jobs (~10%) may be explained by the submission of multi-core jobs that consume more resources than the single core ones. During the 2015 the number of multi core/parallel jobs has enormously increased, showing the results of the effort of adapting scientific applications to the modern CPU architectures.

It is reported in Figure 12 and in Figure 13 the monthly trends about the HEP-SPEC06 hour usage and the number of jobs of the last 3 years respectively. The less increasing trend of the number of jobs can be explained by the increasing popularity of parallel jobs, which use more resources than jobs running on a single job slot. CPU time consumption has increased constantly. The impact of the increasing adoption of multi-threading in jobs, can be noticed in the trends of increasing CPU consumption (CPU normalized hours) from 2008 to date, as plotted in Figure 11 below.

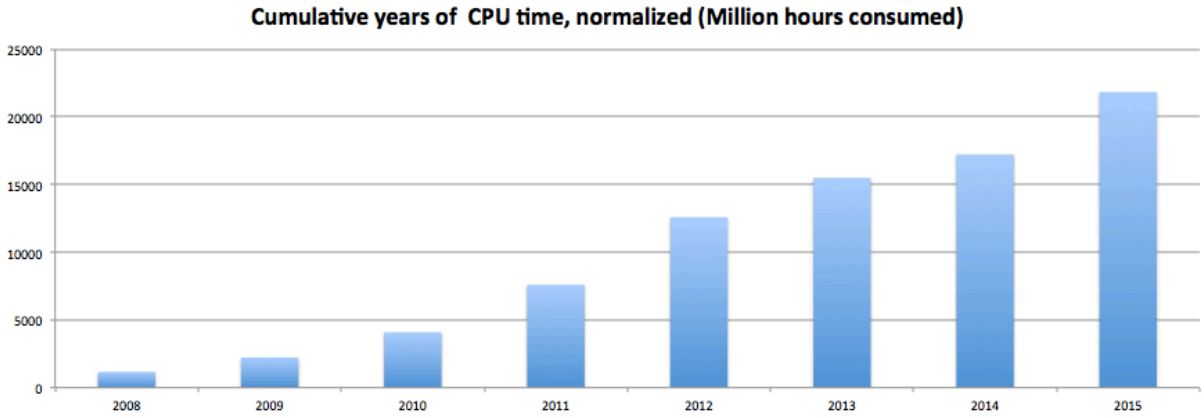

Figure 11 CPU time consumption (normalized time) from 2008 to date.

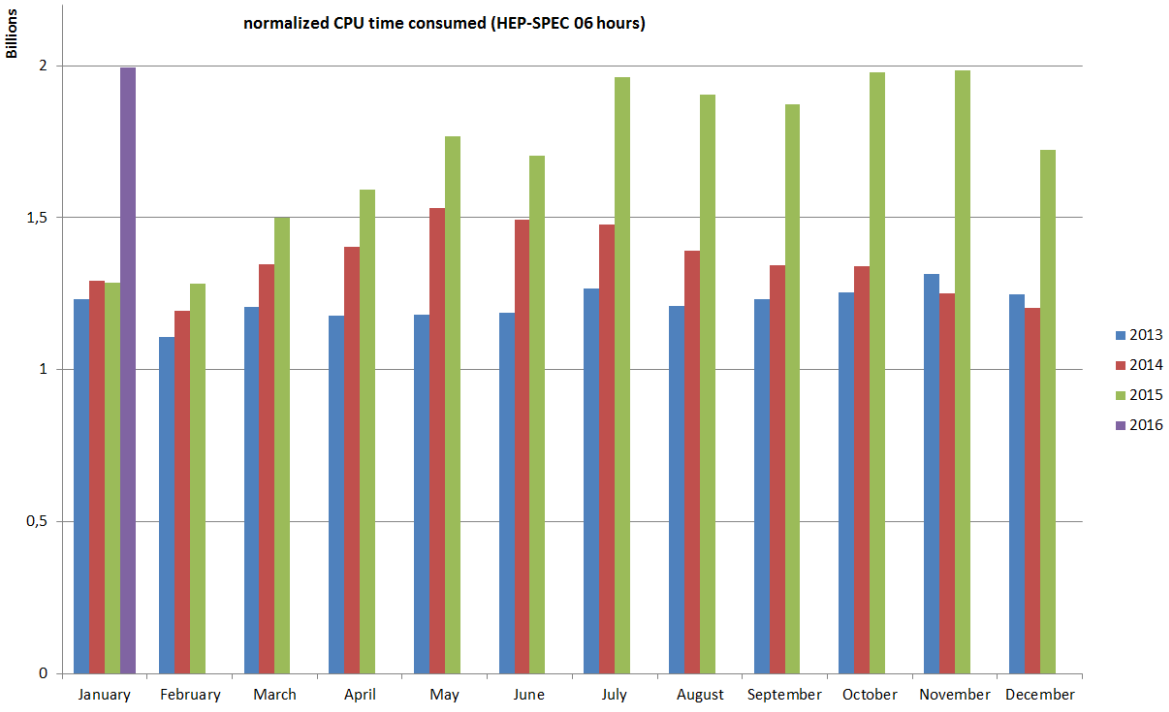

Figure 12 HEP-SPEC 06 Hours monthly usage of the last three years (source: accounting portal).

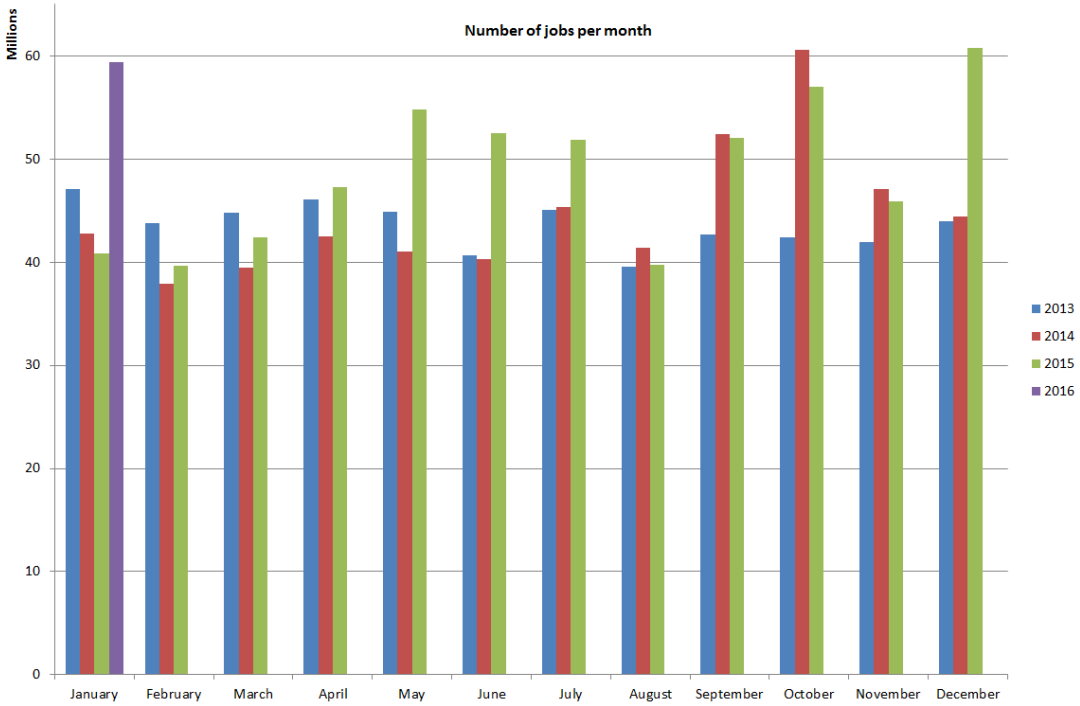

Figure 13 Number of jobs per month of the last three years (source: accounting portal)

The diagrams in Figure 14 and in Figure 15 show the total number of jobs per VO and per Operations Centre respectively, in the period between January 2015 and January 2016.

The usage expressed in HEP-SPEC 06 Hours of CPU wall time across the various resource infrastructures of EGI is plotted in Figure 16, where infrastructures are grouped by operations centre. The diagram also shows the distribution between the four LHC VOs atlas, cms, alice and lhcb (red bars) and the other VOs (blue bars).

The most used infrastructures in absolute terms by the different scientific disciplines (in decreasing order) are: NGI_DE, NGI_UK, NGI_IT, NGI_FRANCE and CERN. Usage distribution naturally reflects availability of installed capacity however the level of multidisciplinary support varies considerably across the infrastructures. Figure 17 plots the distribution of used HEP-SPEC 06 CPU wall clock hours of non-HEP user communities. NGI_IT is the infrastructure with the largest absolute amount of resources used by non-LHC communities with almost 626 Million CPU wall time hours, followed by NGI_DE, NGI_FRANCE, NGI_UK and NGI_TR.

The Figure 18 shows how support of LHC VOs and high-energy physics is dominant in large resource infrastructures, while other disciplines dominate in various countries in Eastern-South Europe. While the LHC VOs altogether account for the largest fraction of resources used in absolute terms, this fraction has been decreasing over during the last years as the services have been increasingly used by new disciplines and projects. Today LHC VO users are less than 50% of the active registered users. The smallest NGIs in terms of number of sites usually support only few VOs, making some NGIs almost discipline-specific. Instead, the larger NGIs that include also large RCs and more regional scientific communities, have the possibility to provide resources to a diverse set of research projects.

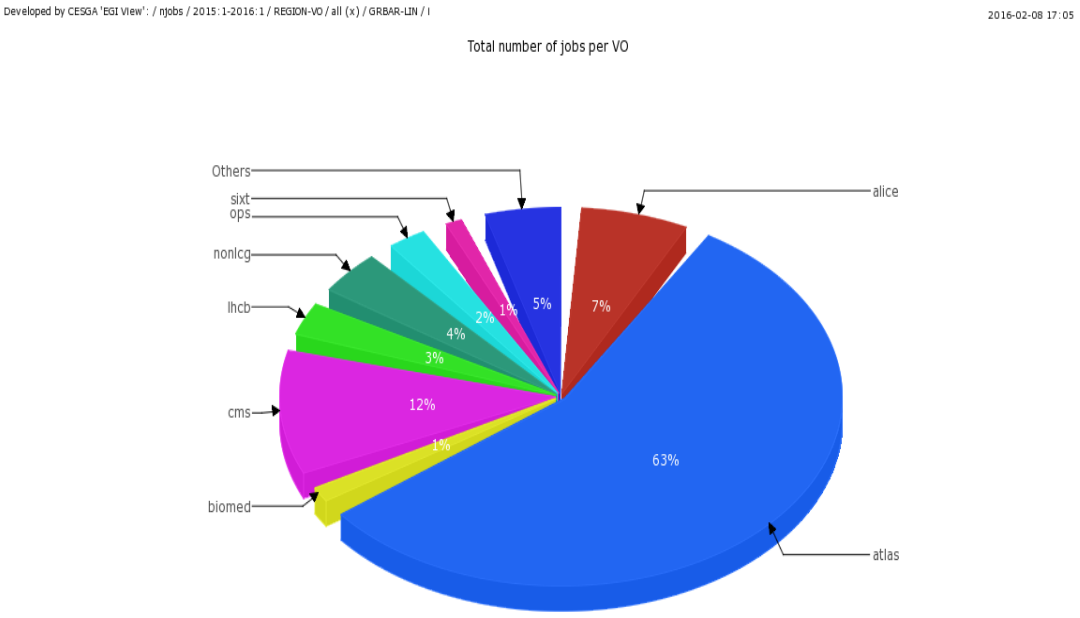

Figure 14 Total number of jobs per VO (Jan 2015 - Jan 2016, source: Accounting Portal)

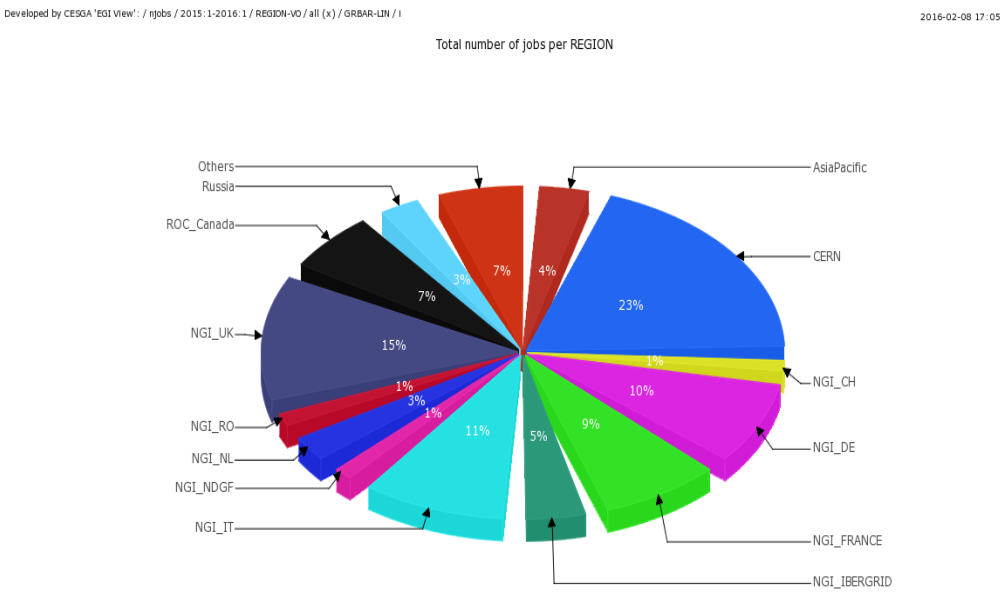

Figure 15 Total number of jobs per NGI/EIRO (Jan 2015 - Jan 2016, source Accounting Portal)

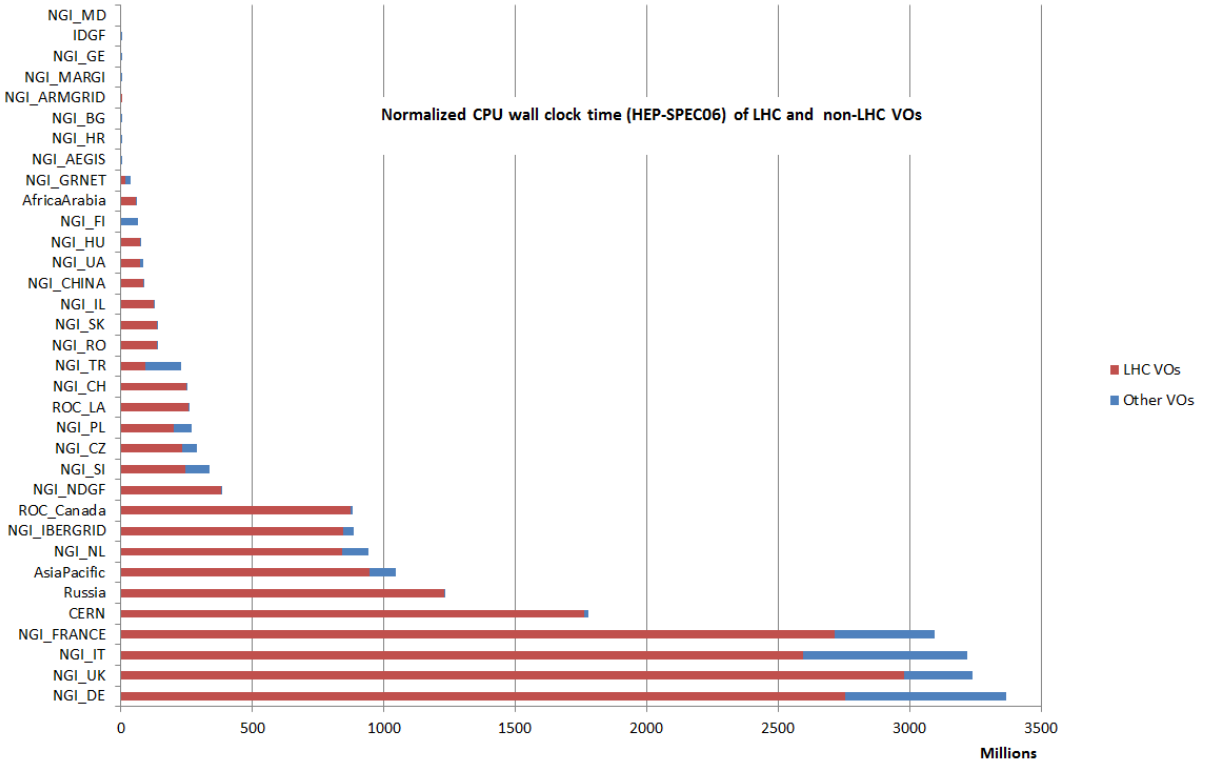

Figure 16 HEP-SPEC 06 Hours from January 2015 to January 2016 (source: accounting portal). LHC usage is displayed in red while the aggregated usage of the rest of VOs is in blue

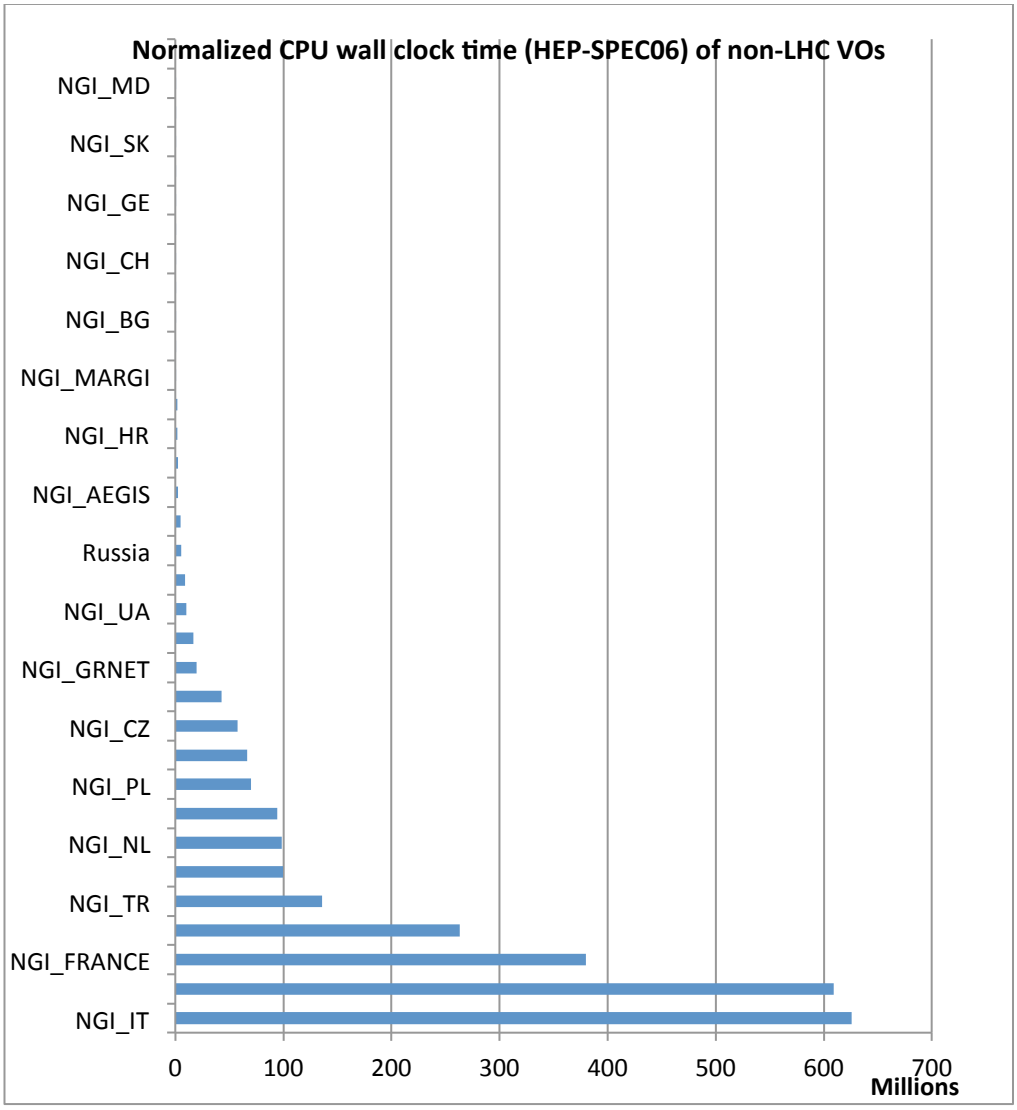

Figure 17 Distribution across EGI Operations Centres of aggregated usage of non-LHC VOs (CPU wall clock time in HEP-SPEC 06 hours) from January 2015 to January 2016 (source: accounting portal).

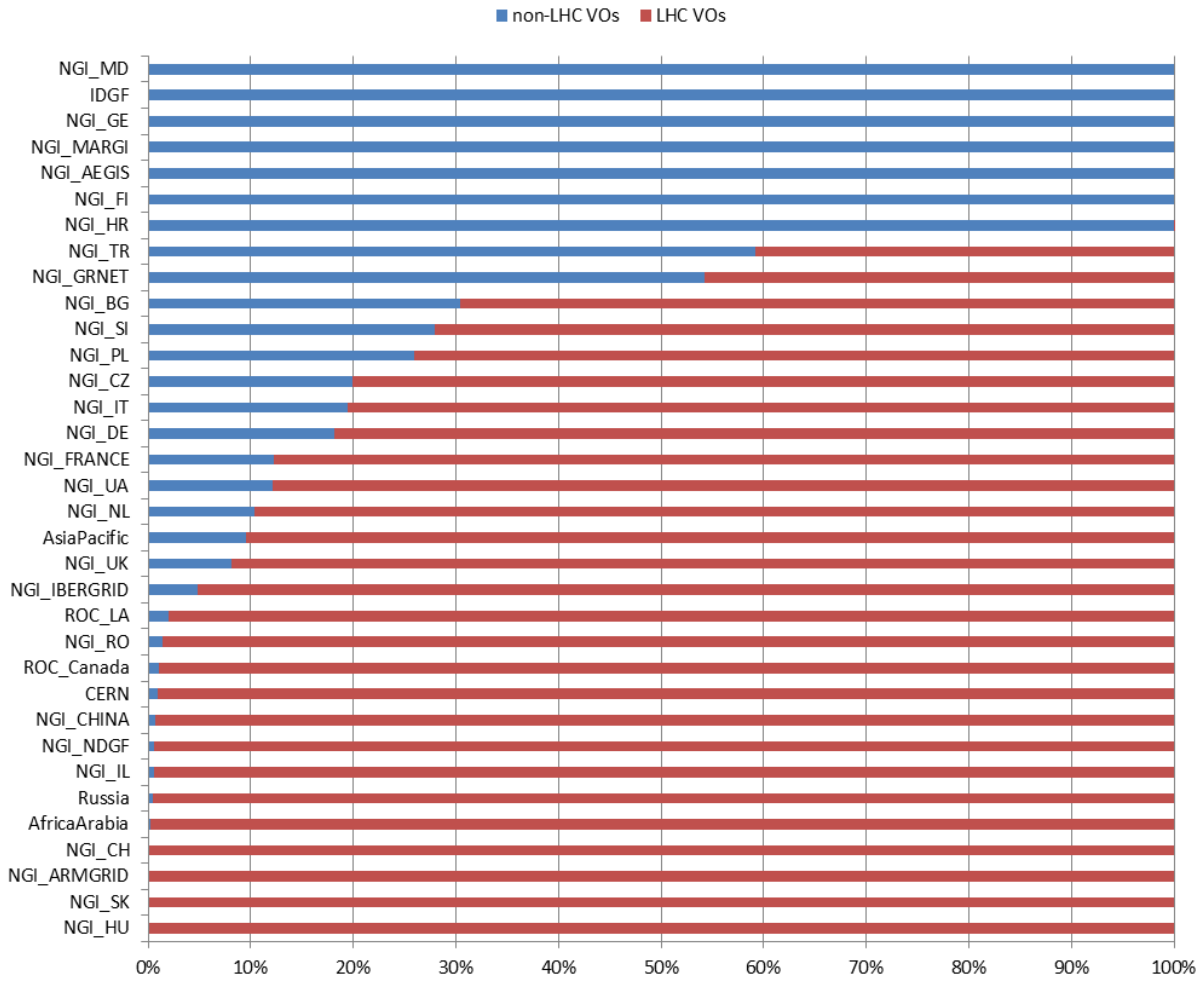

Figure 18 Distribution of resource usage (%) across HEP and non-HEP VOs from January 2015 to January 2016 (source: accounting portal).

As mentioned before, during 2015 accounting of multicore jobs was rolled to production: the not-normalized CPU time consumed by this kind of jobs is reported in Figure 19, while in Figure 20 a comparison of used resources between the single and multi-core jobs is plotted.

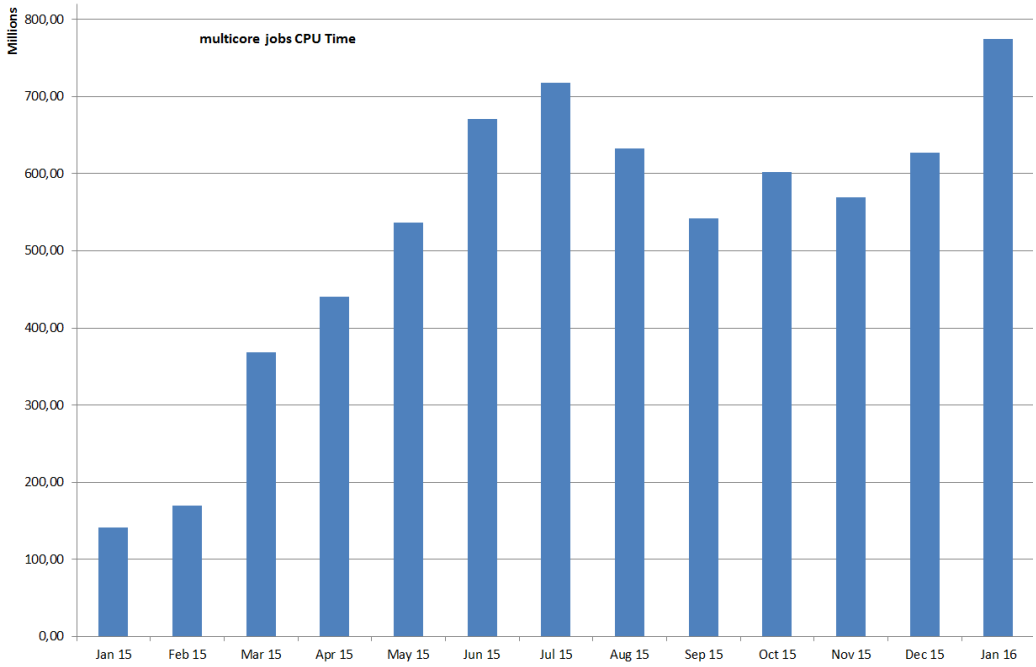

Figure 19 Total (normalised) CPU elapsed time * number of processors. This diagram provides an approximation of the cumulative wall time-equivalent consumption of CPU for multicore jobs.

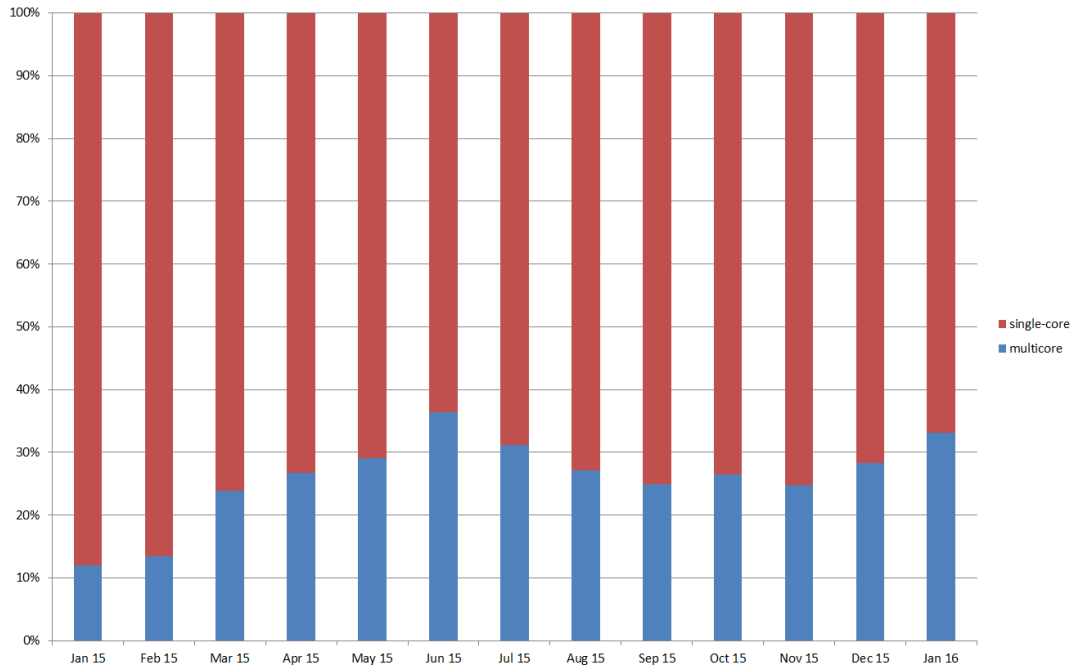

Figure 20 Resource percentage utilization of single and multi-core jobs.

The next graphs shows, considered the total resource usage made by the user belonging to a certain NGI, the percentage of these resources provided by other NGIs.

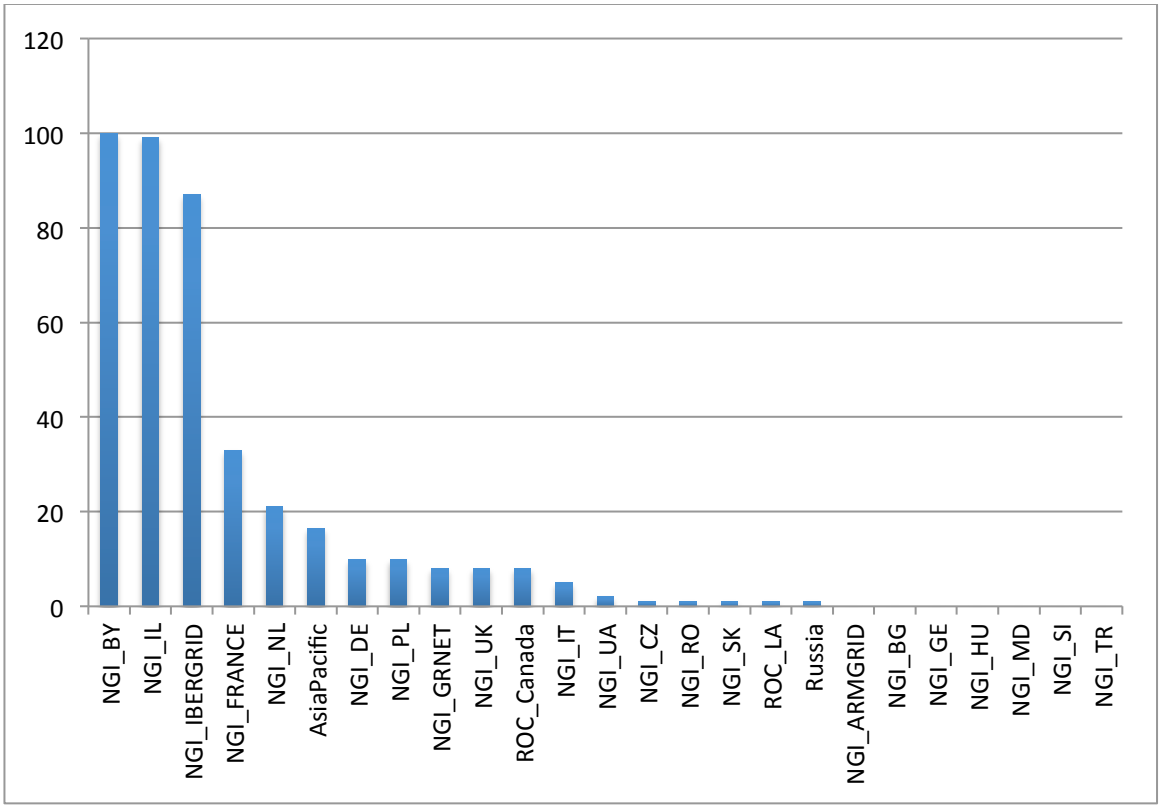

Figure 21 Percentage of the CPU time consumed by national users provided by the resources of other NGIs or EIROs.

One of the main advantages of the resource federations is allowing users to access the unused capacity provided by other countries, in doing so increasing the cost efficiency of the national infrastructures. The figure above shows in which percentage the workload submitted by national users lands on resources provided by resource centres located in different countries, the chart provides a qualitative idea of the positive impact of the EGI federation on the actual research work of the users. This is very visible in relatively small NGIs, such as NGI_BY or NGI_IL, whose users are almost entirely using EGI to access services located abroad.

With regards to cloud, the number of virtual machines instantiated since January 2015 is shown in Figure 22, while Figure 23 plots the percentage of total CPU time consumed in the Federated Cloud across the various providers.

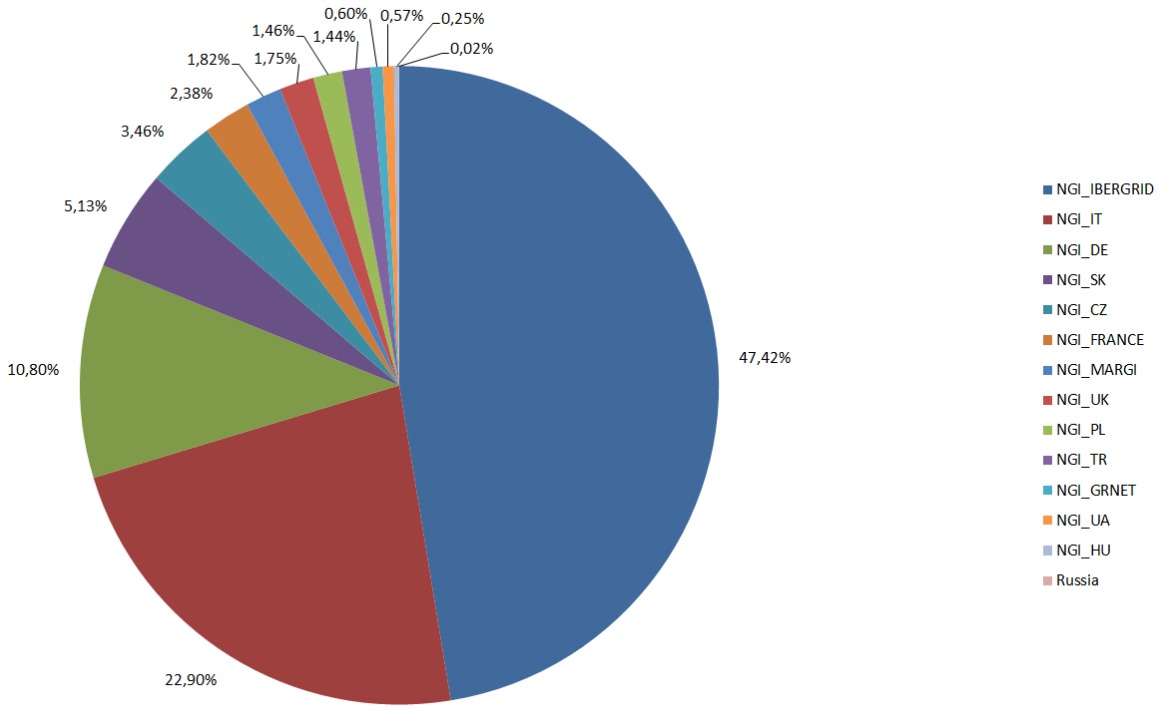

Figure 22 VMs instantiated in the Federated CLOUD per Operations Centre from January 2015 to January 2016 (source: accounting portal)

Table 8 Overall usage of the federated cloud resources between Jan 2015 and Jan 2016 (virtual machines instantiated in the EGI federated cloud, excluding VMs used for monitoring purposes. Source: EGI accounting portal)

Cloud usage during 2015 (relative yearly increase/decrease)	
Total # of VMs instantiated	343,155 (+79%)

The overall number of Virtual Machines instantiated in the federated cloud amounts to 343,155 instances. The platform was launched in production in May 2014; the yearly relative increase compared to the whole of 2014 is +79%. The total amount of CPU wall time hours consumed in 2015 was 2.31 Million. This value does not include the CPU wall time consumed by permanently running VMs, as for this use case, usage is only accounted after termination of the VMs.

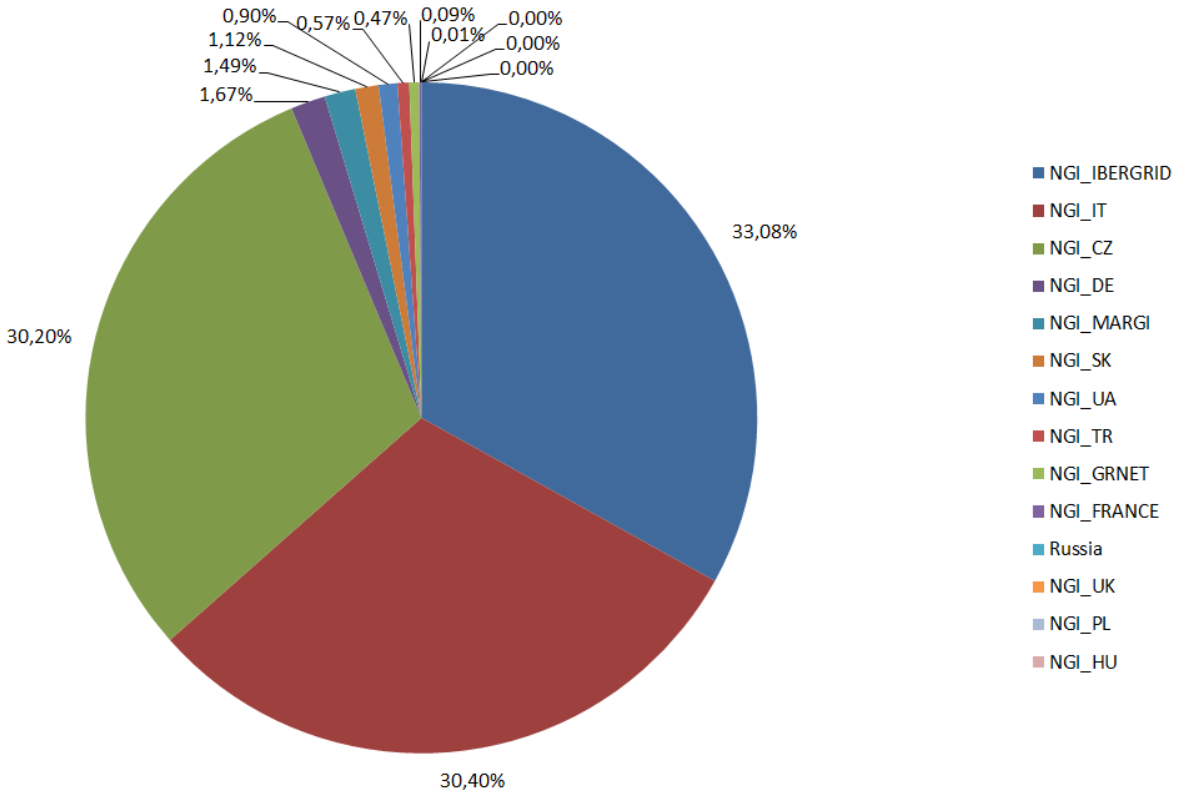

Figure 23 Percentage of total CPU time consumed by cloud services per NGI from January 2015 to January 2016 (source: accounting portal).

2.8 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: Engineering and Technology, Medical and Health Sciences, Natural Sciences, Agricultural Sciences, Social Sciences, Humanities, Support Activities and Others[16]. We should keep in mind that users have different ways of authenticating when accessing services in the distributed infrastructure (e.g. via credentials released by the home organization, personal certificates, and, in the future – where possible – via social network accounts. In addition to this, access can be mediated by platforms or Virtual Research Environments which provide customer-specific tools and services, while relying on baseline e-Infrastructure services. Because of this complexity, the number of active users can only be estimated.

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

The overall number of international and national projects (also known as Virtual Organizations) registered in the Operations Portal[17] at the beginning of February 2016 amounts to 233.

2.8.1 Use of robot certificates

The use of gateways to provide users with a native user-friendly environment to the infrastructure services is increasing. Quite often user portals provide users with the capability of using institutional credentials to authenticate themselves; these credentials are then mapped to robot certificates (often owned by the VO managers). By doing so it is not necessary for a user the request of a personal X.509 certificates and the registration to a VO: this contributes to increase the user friendliness of the platforms. Use of robot certificates is internally accounted for by the portals in compliance to the VO Portal policy. In February 2016 the number of robot certificates embedded in user gateways is 157; robot certificates are used by 51 VOs in total. Almost 11,000 users can potentially use scientific gateways; this is increased by the number of registered users to active VOs, which amounts to be 46246 in February 2016.

The increase in the number of Robot Certificates indicates that users, in particular new user communities, are looking for alternative authentication mechanisms different from the plain X.509 certificates. Within EGI-Engage JRA1 EGI is exploring different authentication technologies and is revising its trust model in order to accommodate the support of different levels of assurance, or to work on a better integration of robot certificates with the production infrastructure.

The diagram in Figure 24 shows the trend in use of robot certificates and VOs since November 2011.

[17] https://wiki.egi.eu/wiki/Scientific_Disciplines

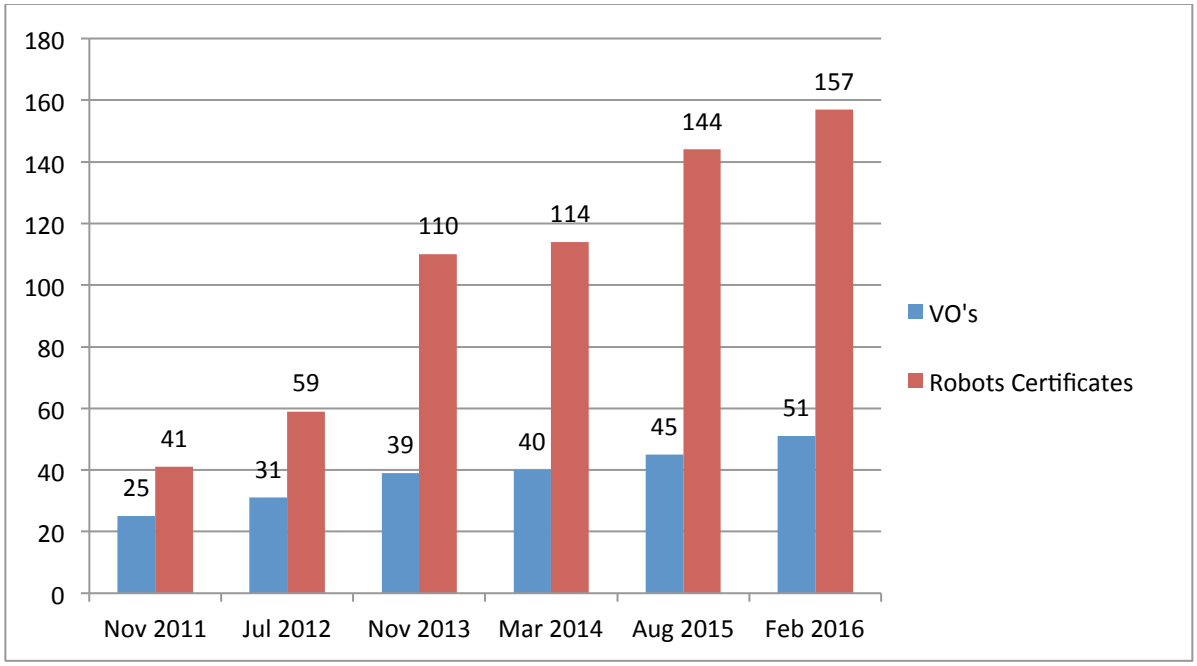

Figure 24 Use of robot certificates and related VO in EGI since 2011.

2.8.2 VOs and user distribution across scientific fields

The distribution of VOs per discipline is illustrated in Figure 25

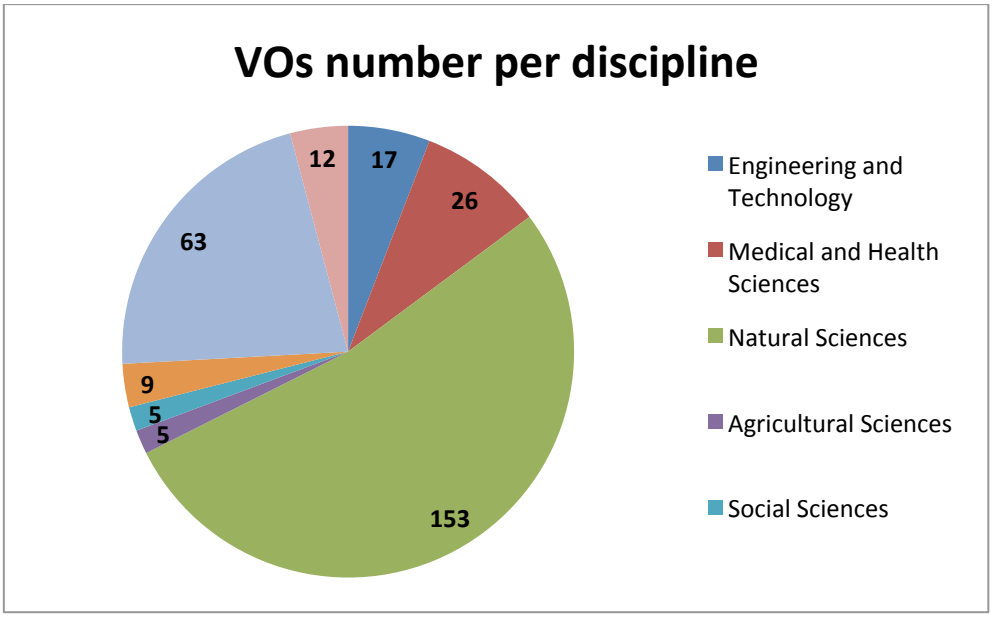

Figure 25 Distribution of number VOs per discipline (February 2016, source: Operations Portal).

The largest discipline in terms of number of registered users is Natural Sciences (65.98%): it is remarkably larger than the other ones because it includes 153 VOs (more than the half of the total VOs). Then there is the Support Activities discipline (9.09%), followed by Medical and Health Science (6,46%) and by Engineering and Technology (6.32%). The complete users distribution is shown in Figure 26.

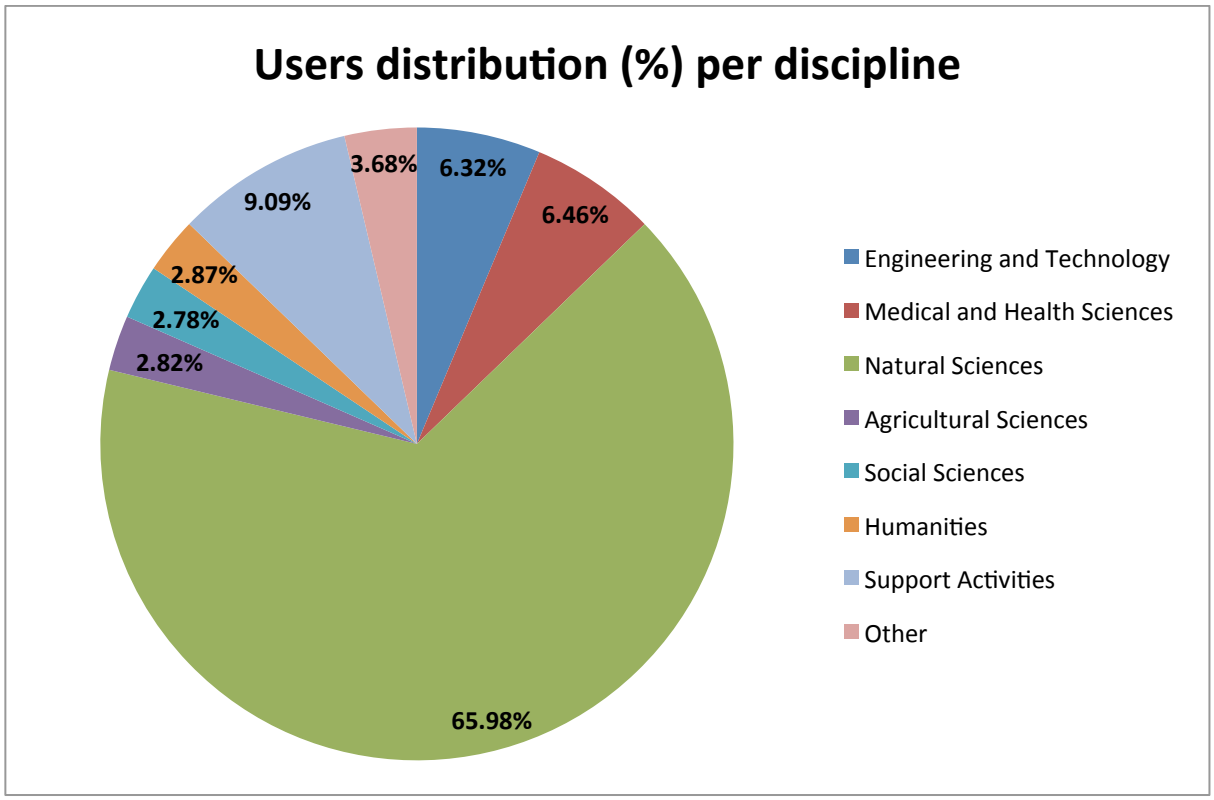

Figure 26 Users distribution per discipline (February 2016, source: Operations Portal). Each VO can be associated to one or more disciplines sub-categories.

2.8.3 Resource utilization per disciplines

Table 9 reports on the increase or decrease in resources usage in 2015 by the 10 most active disciplines in EGI (in terms of HEP-SPEC06 consumed) compared with the data from 2014: the larger disciplines increased the usage in 2015. Computational Chemistry and Medical Imaging scored a significant relative reduction. The medical imaging discipline lead by the BIOMED Virtual Organization, serving multiple independent researchers and small research collaborations (the so called long tail of science) experienced a decrease due to the reduction of the workload produced by the OpenMole application community. The variation is considered to be temporary and related to normal variations in computational needs of the served research communities. On the other hand, the computational chemistry community suffered from service unavailability of the user interfaces operated by the community itself and providing access to the distributed computing platform of EGI supporting it.

Table 9 Increase/decrease of normalised CPU time utilization in 2014 and in 2015 by the 10 most used disciplines (source: accounting portal), ordered by utilization.

DISCIPLINE	Norm. CPU time 2015 compared to 2014
Physics	+28,93%
High Energy Physics	+31,72%
Particle Physics	+27,96%
Nuclear Physics	+43,06%
Space Science	+12,33%
Astrophysics	+12,33%
Astronomy	+9,20%
Medical imaging	-38,45%
Comput. chemistry	-30,94%
Epidemiology	11,97%

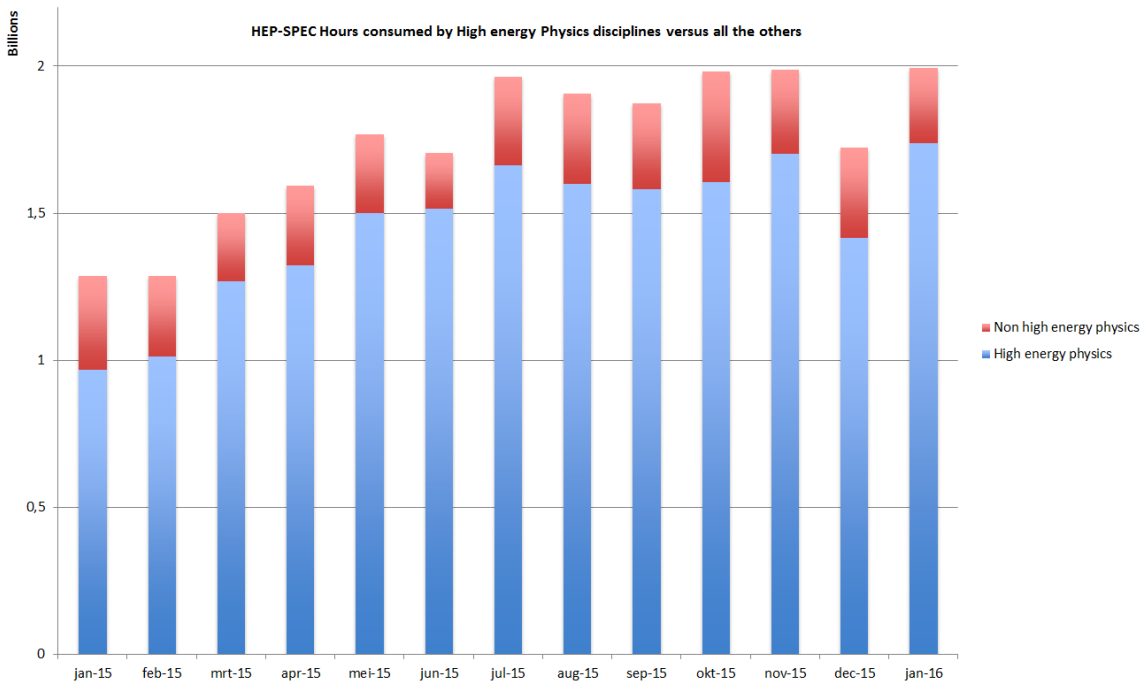

Figure 27 High energy Physics usage compared with all the other disciplines.

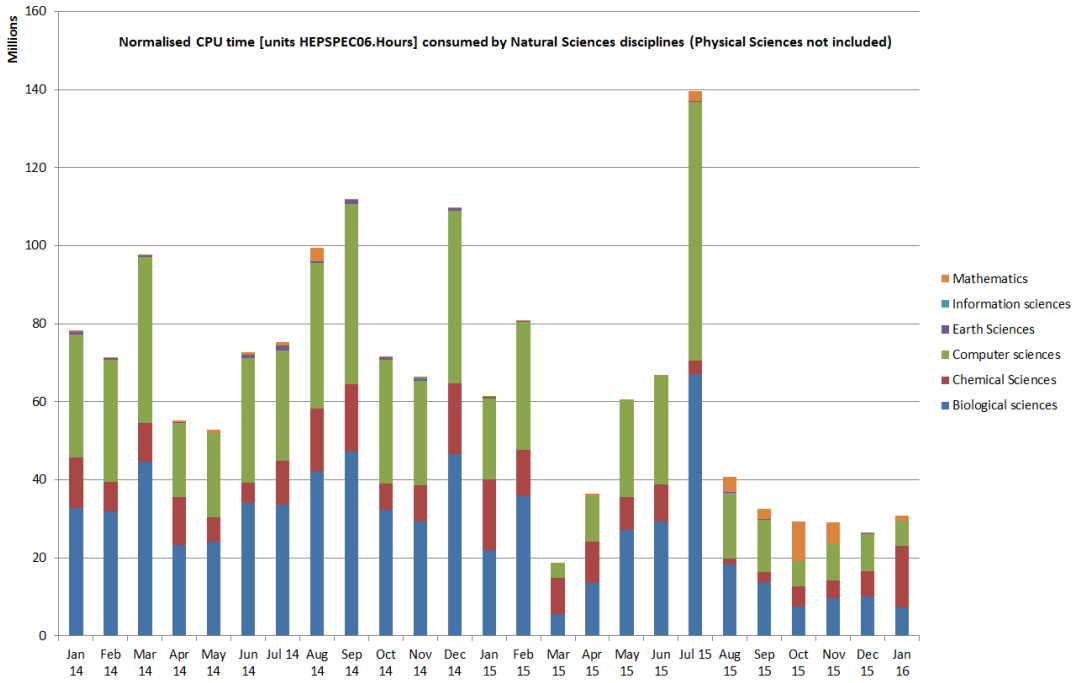

Figure 28 Natural Sciences disciplines resources usage (Physical Sciences excluded).

The figure above shows the usage of resources for the sub-disciplines of the "Natural science" discipline. The "Natural science" discipline has a very broad scope, and in terms of usage is dominated by physics, removing physics the chart shows that the most relevant sub-discipline in terms of capacity consumption is biological sciences.

Also it must be reported that any analysis of accounting data grouped by disciplines can be affected by multi-disciplinary VOs, which reports in more than one discipline.

2.9 Service performances

Services are monitored at three different levels:

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

For each category a different set of service levels and targets are defined and periodically reviewed (see the chapter 3 for details). For each set of service levels various reporting systems are available, and are detailed in the following section. The service levels and targets are formally

defined in the RC Operational Level Agreement, in the RP Operational Level Agreement[18] and EGI.eu Operational Level Agreements[19].

2.9.1 RCs availability and reliability

The quality of HTC 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[20] (SAM). Availability and reliability metrics were defined to quantitatively express the level of functionality delivered by HTC 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. The EGI monthly availability and reliability reports are based on tests (run using the OPS VO) that are sufficiently generic to allow a comparison across all Resource Centres of the infrastructure.

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

Certified Resource Centres need to guarantee 80% minimum availability and 85% minimum reliability for their services (in a distributed environment, workload is distributed across the infrastructure with resilient mechanisms, meaning that the temporary unavailability of a service does not impair the end-user experience, as the workload can be redirected to other service ned-points). The minimum availability and reliability values accepted for a Resource Centre are defined in Operational Level Agreements established with EGI.eu, which is periodically updated.

Increasing the overall performance delivered to users has been an on-going effort since the introduction of service level management.

The trend of the overall EGI RC availability and reliability is shown in the following two diagrams.

[18] https://documents.egi.eu/document/463

[19] https://documents.egi.eu/document/2456

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

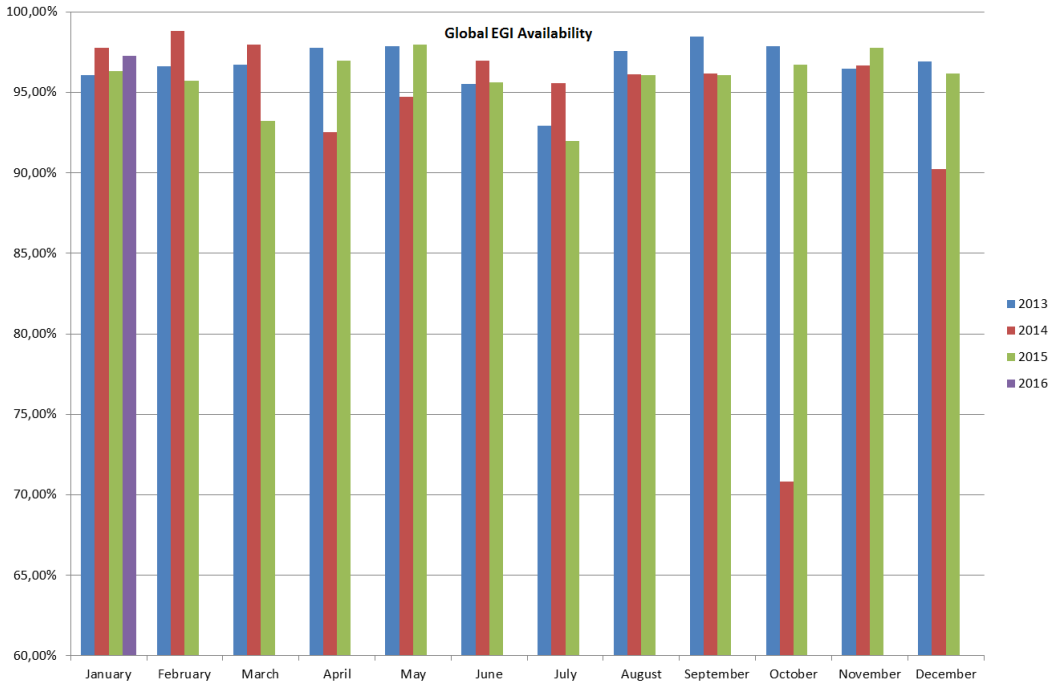

Figure 29 Monthly Availability of resource centres averaged across EGI.

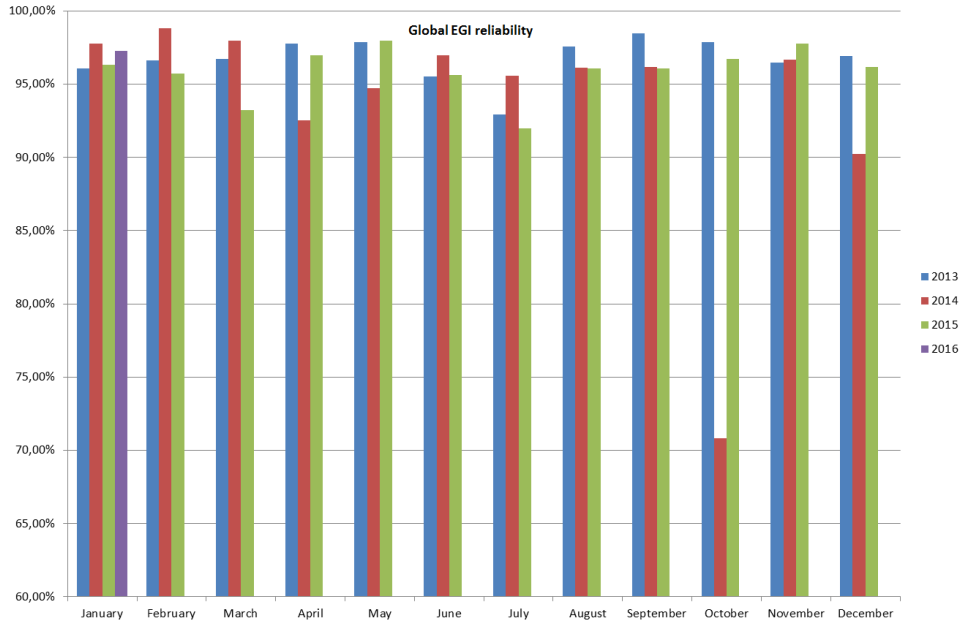

Figure 30 Monthly Reliability of resource centres averaged across EGI, for the last three years.

The overall average availability of the EGI production infrastructure has kept constant and exceeded 95% for almost all the months of 2015. Although it has not improved significantly from

the previous year, average 95% of availability is a good result considering that the highly distributed nature of the infrastructure allows users to use another site if one is not available. Moreover availability is also strictly considering all the site's services, and – for example – if a user is using only the computing services may not be affected by an outage of the storage services.

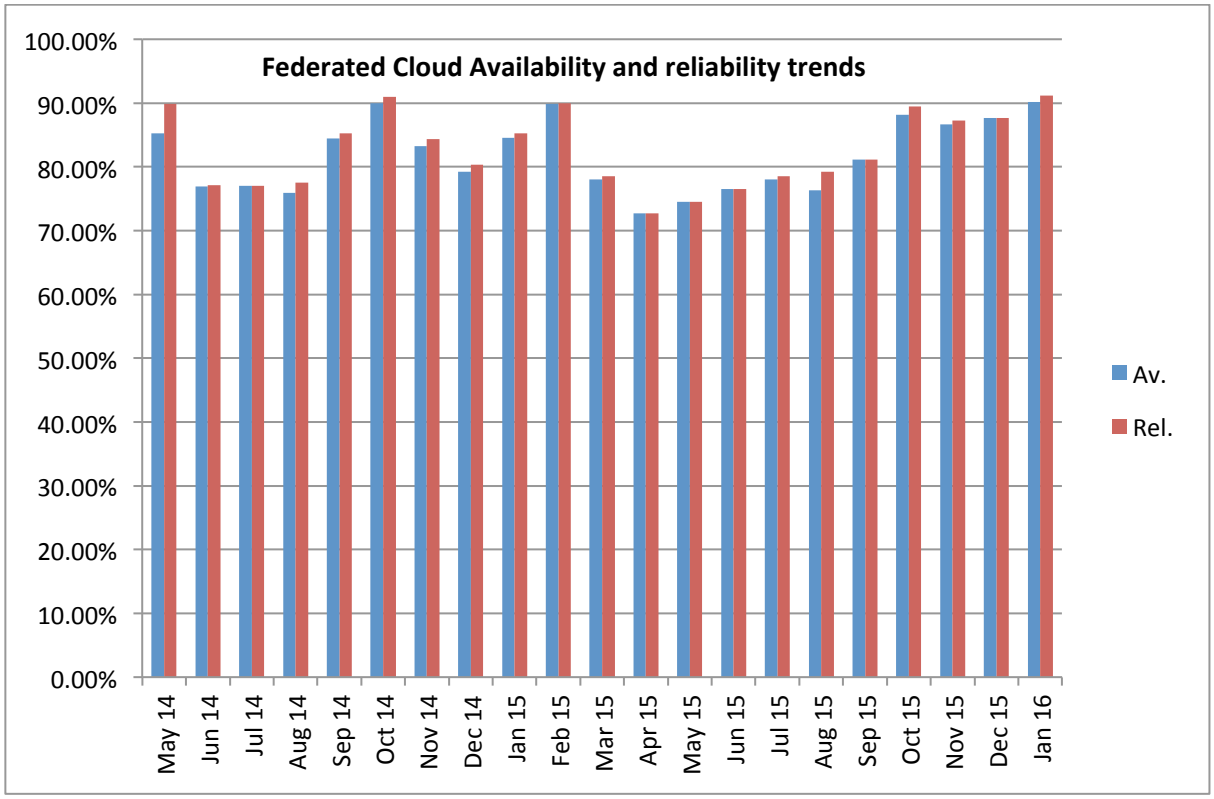

Figure 31 Monthly availability and reliability trends of Federated CLOUD.

Figure 31 shows the availability and reliability trends for the cloud providers: differently from the values computed for the EGI Production Infrastructure, they are not weighted on the capacity deployed by the size, so that small RCs influence the global trends in the same way than the bigger ones; besides the occurrence of problems in small RCs is higher than in the larger ones, so this explains why the average availability and reliability figures of Federated CLOUD are lower than the EGI ones.

As already written, the Federated Cloud Infrastructure started in mid-May 2014. The reports were produced separately from the EGI production infrastructure ones because the Federated Cloud was a test environment during the first part of its life, but in the second half of 2015 there was a general improvement of the quality of the service provided. By now the federated Cloud reached a level of maturity that allows the services to be included in the global EGI availability and reliability computation. As decided In January 2016, EGI federated cloud providers will be subject to the same follow-up operational procedures valid for the HTC RCs, and this is expected to become valid starting as of June 2016.

3 Evolution in the operations coordination

The operations of the EGI production infrastructure are organized at different levels:

- **Infrastructure level**: operations of the federation are coordinated by EGI.eu and the EGI Operations Management Board
- **National level**: operations are organized by the NGI Operations Centres (NOC)
- **Local level**: operations are managed by the data centre staff

This chapter focuses on the operations at e-Infrastructure level, these include:

- Operational processes, procedures, manuals, best practices and policies
- Software provisioning and distribution
- EGI federation services ("core activities and services")

3.1 Operational procedures and processes

Documents are produced by EGI Operations to establish coherent and repeatable procedures for the partners of EGI. While manuals are technical documents that provide guidelines focused on a specific task, procedures are step-by-step descriptions of processes requiring actions from several partners. The purpose of a procedure is to define a workflow. Procedures are approved by the OMB and periodically reviewed. Applicable areas for procedures are:

- Ticket management
- Operations Center Management
- Resource Centre Management
- Availability and monitoring
- Security Incident Handling
- Vulnerability Issue Handling

In the first year of the project, effort concentrated on the development and adaptation of operational procedures and processes for the EGI federated cloud. Pre-existing operational procedures and processes (already adopted for the management of the distributed HTC platform) were adopted as much as possible.

The list below defines the documents includes both new and old documents revised during the last year. The whole collection of operational documents is available on the EGI wiki[21] and it includes: 23 procedures, 12 manuals, 18 HOWTOs, and 10 "Frequently Asked Questions".

[21] https://wiki.egi.eu/wiki/Documentation

Title	New/Updated	Description
Setting up Cloud Resource Centre https://wiki.egi.eu/wiki/MAN10	Updated	It provides very detailed instructions to integrate a cloud infrastructure in the EGI federation. The manual provides information for different cloud middleware stacks supported in the EGI Federated Cloud. Contributions come from both developers of the tools and administrators themselves. The steps to set up a cloud infrastructure in EGI are now well known and straightforward.
Per-User Sub-Proxy https://wiki.egi.eu/wiki/MAN12	New	This manual shows how to set up a per-user sub-proxy (PUSP), which allows identification of the individual users under a common robot certificate. This new feature, defined and developed in the Federated Cloud context, allows a web portal to map a group of users (i.e. VO users) creating a proxy credential from the robot credential. This is fundamental to enable *science gateways*.
Quality verification of monthly availability and reliability statistics https://wiki.egi.eu/wiki/PROC04	Updated	The document describes the process of how to handle justification for poor monthly performance, with the goal of maintaining a given level of quality for the overall EGI infrastructure. The main update to this procedure is represented by a new step in the procedure for communicating with underperforming sites. In case the NGI does not respond to the GGUS ticket in 10 working days, a direct email is sent to the NGI for comments, improving reliability of the communication with the NGI in case of issues through the GGUS ticket, especially with NGIs that are experiencing a frequent personnel turn-over or temporary manpower issues.
Support for CVMFS replication across the EGI Infrastructure https://wiki.egi.eu/wiki/PROC22	New	The procedure describes the process of creating a repository within the EGI CVMFS infrastructure for an EGI VO. This has been entirely defined and tested in the EGI-Engage context.
Production tools release and deployment process	New (still	The procedure describes the process of release and deployment in EGI production

Title	New/Updated	Description
https://wiki.egi.eu/wiki/PROC23	drafted)	infrastructure for Production tools.

3.2 EGI core activities

The EGI core activities and services are fundamental, as they represent the glue that puts together Resource Centres and user communities (Virtual Organizations) implementing the policies between the different partners of EGI. As a consequence, it is very important that they maintain very high levels of availability and reliability. EGI-Engage does not support the operational costs of any of the services operated either at national level or at central level, however, it supports operations coordination, which is responsible for defining the technical specifications, the procurement and the delivery of EGI core activities and services – the EGI service "backbone" providing the glue to federate national services. The EGI core activities and services are partly funded by the EGI participants' yearly fees and partly contributed as in-kind contribution to the federation by the technical partners of EGI.

The specifications of EGI core activities and services are available at:

https://wiki.egi.eu/wiki/EGI_Core_Activities_Bidding#PHASE_II_May_2016-December_2017.

EGI-Engage was responsible for coordinating the procurement of for the period May 2016-December 2017. The bidding process took place during summer 2015. A new service was added, the Application Database, where the EGI Applications Database (AppDB) is a central service that provides:

- Information about software solutions in the form of native software products and virtual appliances, linking the programmers and the scientists who are involved, and the publications derived from the registered solutions.
- The tools for the distribution of the virtual machine images in the cloud sites part of the the federated cloud

Three types of software solutions are offered through the EGI Applications Database:

- Software items, in its classical sense, i.e. applications, tools, utilities, etc.
- Virtual Appliances: composed by one or more pre-configured virtual machine images packaged with an operating system and software application(s).
- Software Appliances: one or more a set pairs of a virtual appliance and a contextualization script. A Contextualization Script (CS) is the script launched on VM boot time and could be used for installing, configuring and preparing software upon boot time on a pre-defined virtual machine image.

The following paragraph provides an overview of the EGI operations services and activities' level targets formally agreed between resource providers, and periodically reported on a monthly basis.

Name	Description	Documentation or service URL
Message Broker Network	The message broker network is a fundamental part of the operations infrastructure ensuring message exchange for monitoring, the operations dashboard and accounting. As such it is a critical infrastructure component whose continuity and high availability configuration must be ensured. The Message Broker Network is part of the EGI Core Infrastructure Platform which is needed to support the running of tools used for the daily operations of EGI.	https://wiki.egi.eu/wiki/Message_brokers
Operations Portal	The Operations Portal provides VO management functions and other capabilities which support the daily operations of EGI. It is a central portal for the operations community that offers a bundle of different capabilities, such as the broadcast tool, VO management facilities, a security dashboard and an operations dashboard that is used to display information about failing monitoring probes and to open tickets to the Resource Centres affected. The dashboard also supports the central grid oversight activities. It is fully interfaced with the EGI Helpdesk and the monitoring system through messaging. It is a critical component as it is used by all EGI Operations Centres to provide support to the respective Resource Centres. The Operations Portal provides tools supporting the daily running of operations of the entire infrastructure: grid oversight, security operations, VO management, broadcast, availability reporting.	http://operations-portal.egi.eu/
Accounting Repository	The Accounting Repository stores user accounting records from various services offered by EGI. It is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI. The EGI Accounting Infrastructure is distributed. At a central level it includes the repositories for the persistent storage of usage records. The central databases are populated through individual usage records published by the Resource Centres, or through the publication of summarised usage records. The Accounting Infrastructure is essential in a service-oriented business model to record usage information.	http://accounting.egi.eu/egi.php
Accounting and Metric	The Accounting Portal provides data accounting views for users, VO Managers, NGI operations and the	http://accounting.egi.eu

Name	Description	Documentation or service URL
Portal	general public. The Accounting Portal is part of the EGI Core Infrastructure Platform which supports the daily operations of EGI. The EGI Accounting Infrastructure is distributed. At a central level it includes the repositories for the persistent storage of usage records. The central databases are populated through individual usage records published by the Resource Centres, or through the publication of summarised usage records. The Accounting Infrastructure is essential in a service-oriented business model to record usage information. The Metrics Portal aggregates metrics from the EGI Infrastructure from activity leaders and NGI managers in order to quantify and track the infrastructure evolution.	
SAM central services	The Service is part of the EGI Core Infrastructure Platform which supports the daily operations of EGI. Central systems are needed for accessing and archiving infrastructure monitoring results of the services provided at many levels (Resource Centres, NGIs and EGI.EU), for the generation of service level reports, and for the central monitoring of EGI.eu operational tools and other central monitoring needs. The system is currently going to be moved from the old distributed MyEGI to the new (central) ARGO infrastructure.	http://argo.egi.eu/ https://wiki.egi.eu/wiki/SAM_Instances
Monitoring central services	Monitoring Central Services is supporting monitoring of activities to be conducted centrally, like monitoring of e.g. UserDN publishing in accounting records, GLUE information validation, software versions of deployed middleware, security incidents and weaknesses and EGI.eu technical services. Central Monitoring Services is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI.	
Security monitoring and related support tools	Security monitoring and related support tools are part of the EGI Core Infrastructure Platform which supports the daily security operations of EGI. EGI is an interconnected federation where a single vulnerable place may have a huge impact on the whole infrastructure. In order to recognise the risks and to address potential vulnerabilities in a timely manner, the EGI Security Monitoring provides an oversight of the	https://wiki.egi.eu/wiki/EGI_CSIRT:SMG

Name	Description	Documentation or service URL
	infrastructure from the security standpoint. Also, sites connected to EGI differ significantly in the level of security and detecting weaknesses exposed by the sites allows the EGI security operations to contact the sites before the issue leads to an incident. Information produced by security monitoring is also important during assessment of new risks and vulnerabilities since it enables to identify the scope and impact of a potential security incident.	
Service registry (GOCDB)	Service Registry (GOCDB) is a central registry to record information about different entities such as the Operations Centres, the Resource Centres, service endpoints and the contact information and roles of people responsible for operations at different levels. GOCDB is a source of information for many other operational tools, such as the broadcast tool, the Aggregated Topology Provider, the Accounting Portal, etc. GOCDB is part of the EGI Core Infrastructure Platform, which supports the daily operations of EGI.	http://goc.egi.eu/
Catchall services	Catch-All services are auxiliary services needed by the Core Infrastructure Platform and by various operational activities of EGI. Auxiliary services and activities are needed for the good running of Infrastructure Services. Examples of such services are VOMS service and VO membership management for infrastructural VOs (DTEAM), the provisioning of middleware services needed by the monitoring infrastructure (e.g. top-BDII and WMS), and catch-all services for emerging user communities.	https://wiki.egi.eu/wiki/Catch_All_Grid_Core_Services
Operations support	Operations support is auxiliary service needed by the Core Infrastructure Platform and by various operational activities of EGI. Auxiliary activities are needed for the good running of Infrastructure Services. Examples of such are activities for service level management, service level reporting, service management in general and central technical.	https://wiki.egi.eu/wiki/EGI_Infrastructure_operations_oversight
Security coordination	Central coordination of the security activities ensures that policies, operational security, and maintenance are compatible amongst all partners, improving integrity and availability and lowering access barriers for use of	https://wiki.egi.eu/wiki/Security

Name	Description	Documentation or service URL
	the infrastructure.	
Acceptance criteria	The Acceptance Criteria are the functional and non-functional requirements that a product must fulfil to be released in UMD, these include generic requirement applicable to every product, and specific requirements applicable to the capabilities supported by a component.	https://wiki.egi.eu/wiki/Software_Provisioning_Process
Collaboration tools/IT support	Collaborations tools are services needed by the EGI back-office and supporting EGI collaboration.	https://wiki.egi.eu/wiki/EGI_Collaboration_tools
Staged Rollout	The Staged Rollout is an activity by which certified updates of the supported middleware are first tested by Early Adopter (EA) sites before being made available to all sites through the production repositories. This procedure permits to test an update in a production environment that exposes the product to more heterogeneous use cases than the certification and verification phase. This allows the discovery of potential issues and potentially to add mitigation information to the UMD release notes.	https://wiki.egi.eu/wiki/Staged_Rollout
Software provisioning infrastructure	The software-provisioning infrastructure provides the technical tools to support the UMD release process from pulling packages from the developers' repositories to the build of a release.	https://wiki.egi.eu/wiki/EGI_Software_Component_Delivery
Incident management helpdesk	Incident Management (Helpdesk) is the central helpdesk provides a single interface for support. The central system is interfaced to a variety of other ticketing systems at the NGI level in order to allow a bi-directional exchange of tickets. GGUS is part of the EGI Collaboration Platform and is needed to support users and infrastructure operators.	http://helpdesk.egi.eu
1st and 2nd level support (core platform, community platform)	First level support is responsible for ticket triage and assignment. This activity is also responsible for the coordination with teams responsible for 2nd level and 3rd level support. Software-related tickets that reach the second level of support are analysed and if necessary are forwarded to 3rd line support units only when there are clear indications of a defect (in software, documentation, etc.).	

Name	Description	Documentation or service URL
AppDB	The EGI Applications Database (AppDB) is a central service that stores and provides to the public, information about software solutions in the form of native software products and/or virtual appliances, the programmers and the scientists who are involved, and publications derived from the registered solutions.	https://appdb.egi.eu/
e-Grant	e-GRANT is a tool supporting Resource Allocation process. It allows researchers to request an amount of compute and storage resources, or FedCloud resources, for a given amount of time. e-GRANT handles all activities involved in RA Process which leads to SLA signing.	https://e-grant.egi.eu/slaneg/auth

During EGI-Engage the provision of the core activities, has been regular and – with minor deviation – in the boundaries set by dedicated OLA signed by the service providers.

3.3 UMD software provisioning

The Software Provisioning infrastructure provides the technical tools to support the UMD release process from pulling packages from the developers' repositories to the build of a release. The main goals of the overall Software Provisioning process are:

- Distributing the software provided by the Technical Providers (i.e. development teams) through a central repository
- Verify that the software fulfils a given set of Quality Criteria
- Deploy the software into the infrastructure in a controlled way, so that software is installed on sites only if it has been tested in a real production context (Early Adoption).

The UMD Software Provisioning activity is made of several components:

- **Software Provisioning Process**, made of 3 sub-processes
  - **Software Delivery**, when Technology Providers (i.e. the development/product teams) submit new software releases; this is made by email or GGUS ticket. Software delivery is performed using one of the three different user interfaces available, i.e. a web form, e-mailing and a web service interface, that create tickets including all the necessary information about the software delivered in order to be processed. GGUS forwards the tickets to RT creating one RT ticket per Product per Platform and Architecture (PPA)
  - **Software Assessment**, consisting in

- **Quality Assurance**, which assures that the software fulfils the Quality Criteria to be released in UMD; these include generic requirements applicable to every product, and specific requirements applicable to the capabilities supported by a component. During the last year, a new version of the Quality Criteria has been produced[22]
- **Staged Rollout**, which is a procedure by which certified updates of the supported middleware are first tested by Early Adopter (EA) resource centres before being made available to all sites through the production repositories. This procedure permits to test an update in a production environment that exposes the product to more heterogeneous use cases than the certification and verification phase. This allows the discovery of potential issues and the addition of corresponding mitigation information to the UMD release notes.

- **Reporting**, which is about informing TPs about the outcome of the Software Provisioning Process

- **UMD Release Process**, collecting tested Products per Platform and Architecture (PPAs) into UMD Releases.

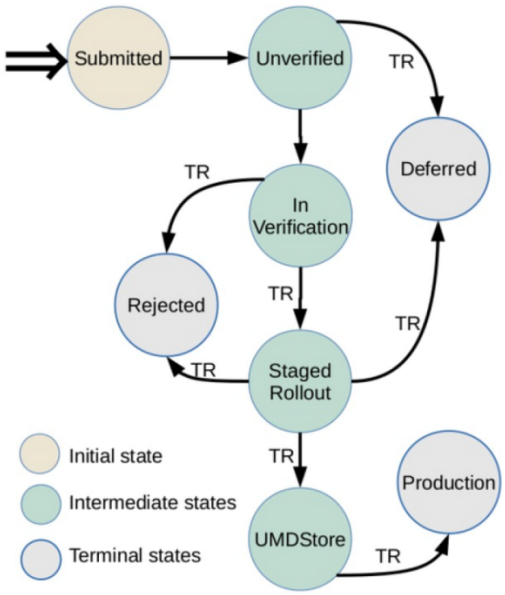

Figure 32 Status of products through the Software Provisioning Process

[22] http://egi-qc.github.io/

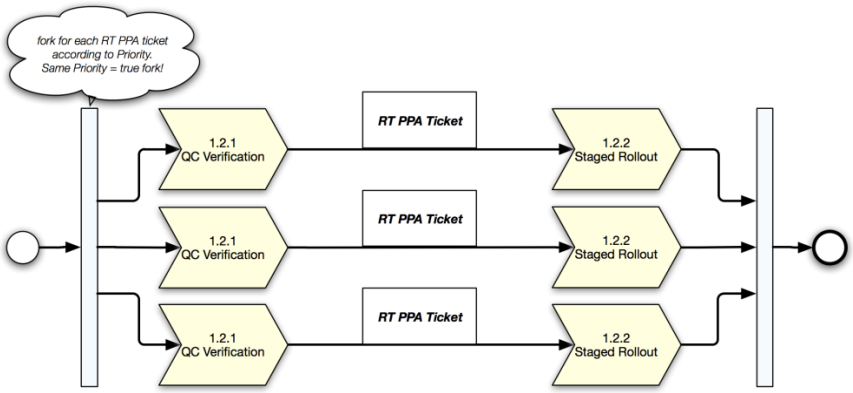

Figure 33 Software assessment

The Software Provisioning infrastructure is composed by several components. The most important are:

- **RT** (Request Tracker) tracks the status of the product in the software provisioning process, for a given release of a given product;
- **Repository Back-End** automates the movement of packages between repositories, validating the individual product releases submissions
- **Composer**, a web-based interface for bundling versioned software products that have successfully passed the UMD verification process, into a robust UMD release ready to be deployed
- **Web frontend** publishing the information about UMD releases
- **Repositories** to be maintained for every operating system and major release supported; they are:
    - **Untested**: contains the packages to be installed during the verification
    - **Testing**: contains the packages to be installed during staged rollout
    - **Base**: contains the packages released in the first major release
    - **Update**: contains the packages released in the update releases
    - **Release Candidate**: it is generated before a UMD release, to simulate the production repositories after the UMD release under preparation. This is used to test the installability of the newly released components, as well as the products already in production.

The Software Provisioning infrastructure supports multiple operating system (EL-based, Debian-based) and major releases. At the moment, the UMD4 structure provides support to CentOS7, SL6, and Ubuntu.

The infrastructure provides also a *"Preview" repository* where products are quickly released without verification; this is not an official UMD repository, but it follows the same procedures and has the same features.

The UMD Release Team (URT) has been working during the last year on releasing the middleware distribution according to the criteria of the Software Provisioning Process developed in the EGI context.

The documentation and the procedures used by the URT have been modified to improve the performances and the reliability of the whole process. In particular, a new guide has been written to make the whole process more transparent, to provide tracking of the packages and feedback during the followed steps. The provisioning procedure is organized as a process compatible with FitSM standard[23]. Moreover, several optimizations to the release workflow allow now to make the release time more predictable (about 2 months) and more reliable (no products out of the UMD radar).

At the moment, EGI is supporting two different major releases of UMD:

- **UMD3**, supporting two EL-based platforms (Scientific Linux 5 and Scientific Linux 6) and Debian;
- **UMD4**, supporting two EL-based platforms (CentOS7, Scientific Linux 6) and Ubuntu

The choice of the new platforms for UMD4 has been driven by:

- Almost full backward compatibility with previous releases (CentOS with SL, Ubuntu with Debian)
- Explicit preference of the Resource Centres in using CentOS/Ubuntu instead of SL/Debian (survey presented on May 2015).

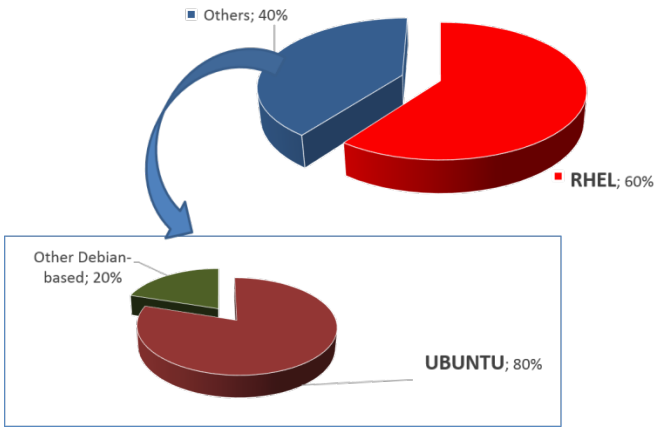

Figure 34 Use CentOS and Ubuntu in the EGI Federated Cloud

A new release of the UMD has been released in January 2016 (UMD 4.0.0). The first release focuses on the adoption of CentOS7. The UMD4 release will soon host also the products developed in the EGI Federated Cloud and specifically in EGI-Engage JRA1 (Cloud Management Framework) for Ubuntu 14.04 LTS.. Considering that SL5 is under the decommissioning phase (the

[23] http://fitsm.itemo.org/

deadline is April 2016), and the SL6 porting from UMD3 to UMD4 is currently ongoing, the decommissioning of UMD3 can be planned at the beginning of PY2 in favor of UMD4.

During 2015 there have been 9 releases of UMD3, 4 of which are minor releases and 5 revisions or fixes.

Table 10 Products updated in UMD3 during 2015

Capability	Product/Technology
Compute, Job Execution, Job Scheduling	CREAM, Globus GRAM5, QCG-Computing, UNICORE TSI, UNICORE/X, ARC, WMS
Accounting	APEL
Attribute authority	VOMS Server, UNICORE XUUDB
Authentication	Globus GSI, UNICORE-Gateway
Authorization	ARGUS-PAP
Client tools	GFAL2 utils, VOMS clients
Credential management	MyProxy, ProxyRenewal
Data access	DAVIX
File Access, File Transfer, Storage Management	StoRM, dpm-xroot, XRootD, CVMFS, dCache, DPM/LFC, Frontier SQUID, Globus GRIDFTP, XROOTD
File Transfer Scheduling	FTS3
Information Discovery	Globus InfoProviderService, UNICORE Registry
Other	BLAH, CGSI-gSOAP, CREAM TORQUE module, CREAM GE module, DMLITE, GFAL2, GFAL2-python, SRM-ifce, classads-libs, edg-mkgridmap, fetch-crl, ARC Nagios probes

3.3.1 User software distribution

While UMD aims at distributing the “middleware” to the Resource Centres, other ways are necessary to distribute user applications: in these cases the life cycle management of the application must be decoupled from the operating system and the middleware as much as possible.

CVMFS (CERNVM File System) fits perfectly to the case and its deployment was improved in 2015. CVMFS is a network file system based on HTTP and optimized to deliver experiment software in a fast, scalable, and reliable way. Files and file metadata are downloaded on demand and locally cached, without interfering with the base system.

Several VOs are asking for migrating to CVMFS to distribute their software. EGI has formalized a procedure to drive the Virtual Organizations through setting up CVMFS for their software in the EGI infrastructure, making the software automatically available at the Resource Centres by means of the preinstalled CVMFS clients.

Effort has been spent to ensure interoperability between the CVMFS services provided by EGI and OSG: as anticipated, a procedure[24] has been set, and recently refined, to ensure that VOs managing the CVMFS area are supported by Resource Centres in both OSG in the United States and EGI.

3.3.2 Virtual Appliance distribution and VA Endorsement

The AppDB acts basically as a catalogue of virtual appliances (VA): for each VA, it maintains a set of metadata, among which a description, an identifier, and the URL of the VA itself, which is not stored on the AppDB itself. Versioning of the VAs is supported as well. After publishing a new appliance, or a new version of an existing appliance, everybody is able to download the instance.

The VO manager manages a "VO image list"; he can add a VA to his VO image list. Then the members of a VO can run the images of their VO image lists on the Resource Centres supporting their VO; this means that the VO manager decides which images can be run on the EGI Federated Cloud by the VO members by simply listing the allowed images in the VO image list. This is the reason why the act of adding a VA to a VO image list is called **endorsement**: the VO manager takes the responsibility for stating that the image is "trusted" and can be used in the context of his/her VO.

The images available in AppDB can be classified in two types:

- EGI images, which are general purpose images, based on broadly used OSes
- VO-specific image, which are VO specific images, available to a specific VO and customized for specific purposes

Procedures are available that can assure that a given virtual appliance published in AppDB, under control of a given VO manager/endorser, is well-configured, secure and up-to-date. A checklist has been drafted to schedule and execute the maintenance of the appliances; the work done so far is available on wiki[25].

In order to make the VAs of a specific VO list available at the Resource Centres supporting the VO, two tools are used: vmcatcher[26] and vmcaster. In particular, vmcaster is a tool for managing and

[24] https://wiki.egi.eu/wiki/PROC20

[25] https://wiki.egi.eu/wiki/Virtual_Machine_Image_Endorsement

[26] Documentation about vmcatcher/caster: https://github.com/hepix-virtualisation/vmcatcher

https://github.com/hepix-virtualisation/vmcaster

updating published virtual machines image lists on the AppDB side. On the other hand, vmcatcher is the corresponding counterpart on the Resource Centre side: for each supported VO, it downloads the corresponding image list, and synchronizes locally a fresh copy of each image; if the list is updated, the vmcatcher takes care of reflecting the modifications locally (adding new images, deleting old images, overwriting old versions with new ones). In fact, this allows the endorser of the images to decide which images should be available at the Resource Centres in a very simple way and centrally for the whole federation.

At the time of writing, the Application Database includes 83 virtual appliances and 10 software appliances. A virtual appliance is a virtual machine image that is used to instantiate a virtual machine, a software appliance is a tuple of virtual appliance plus the contextualization scripts to install and configure specific software on the virtual machine after the instantiation. The main advantage of the software appliance is that a user is not required to create a virtual machine image and register it in AppDB, but they can use an existing virtual appliance and have the needed software installed and configured using the contextualisation script.

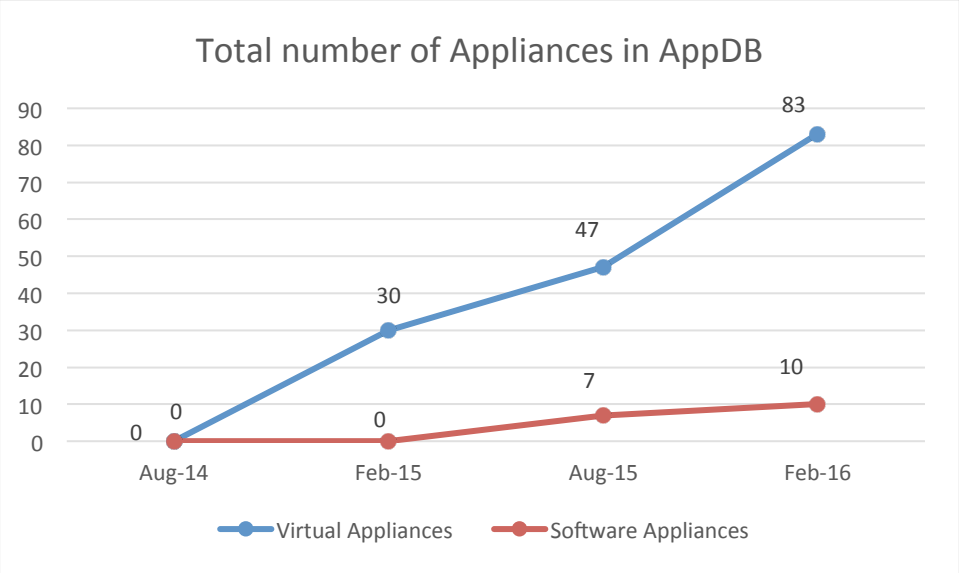

Figure 35 total number of appliances registered in AppDB, trend in the last 18 months

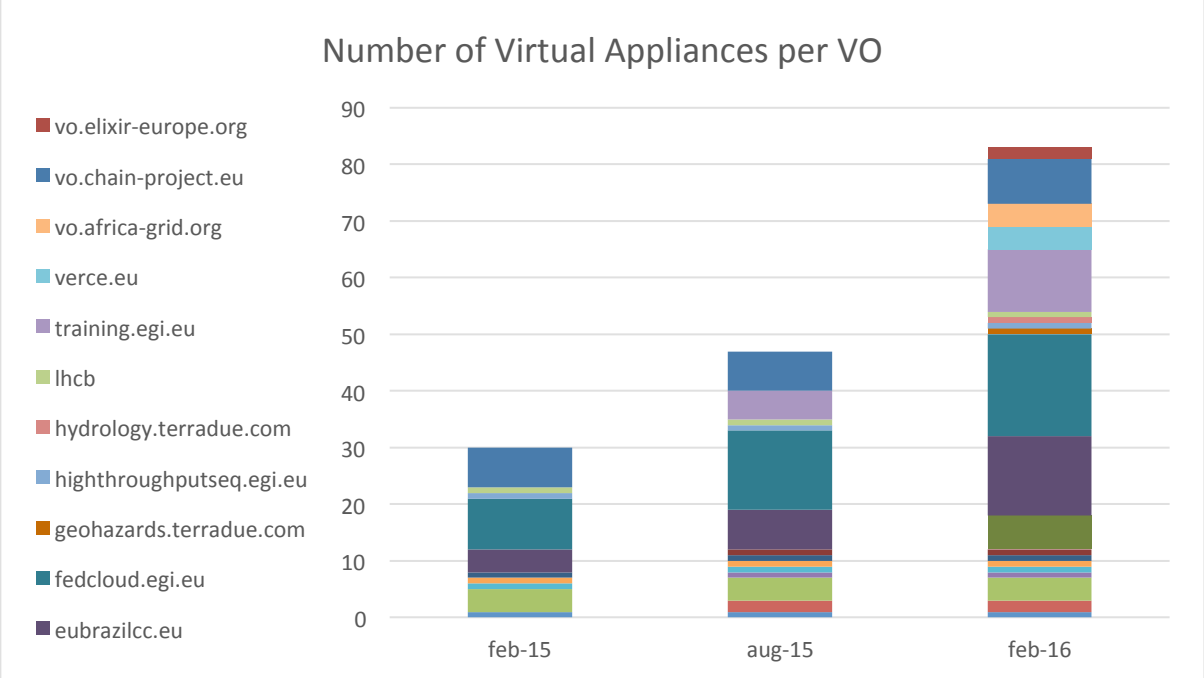

Figure 36 Number of appliances endorsed by VOs

These two graphs provide a general positive trend in the total number of Virtual Appliances and Software Appliances stored by Virtual Organizations, especially in the last 6 months.

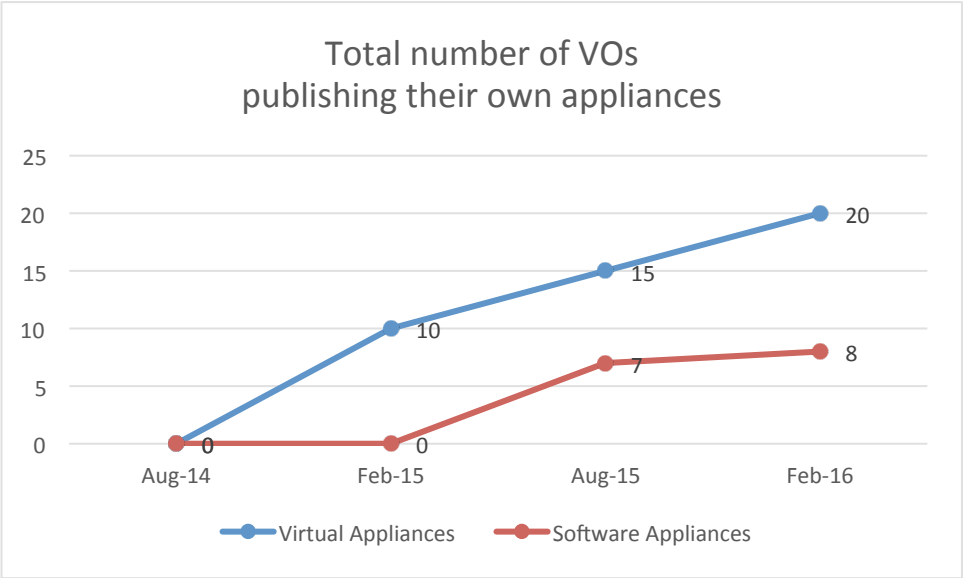

Figure 37 number of VOs using AppDB to publish their own appliances

The number of VOs publishing appliances on AppDB doubled during the course of the last year according to the linear increase plotted in the figure above.

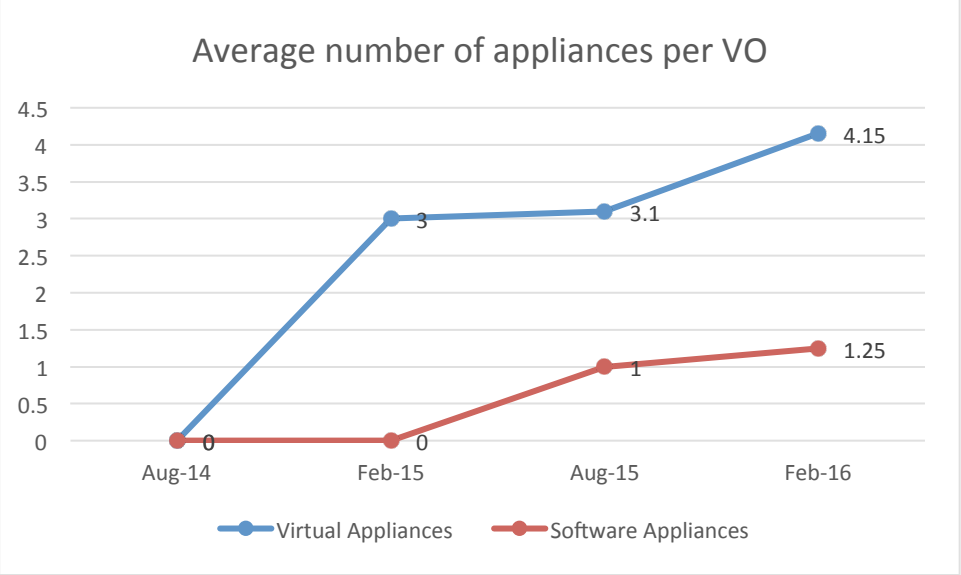

Figure 38 Average number of virtual appliances published by a VO in AppDB

Taking into consideration only the VOs publishing appliances, the number of appliances published by a single VO increased on average by +25% for both software and virtual appliances in 6 months).

3.4 IT service management

During the first year of the project, WP5 has been working towards improvement of EGI service delivery and quality of its services. The ITIL/ISO20k-based standard FitSM was chosen as reference standard, being it a community standard, developed in the context of an EC project, and lightweight for facilitating service management in IT service provision, including federated scenarios. The main goals of FitSM are:

- Create a clear, pragmatic, lightweight and achievable standard that allows for effective IT service management (ITSM).
- Offer a version of ITSM that can cope with federated environments, which often lack the hierarchy and level of control seen in other situations.
- Provide a baseline level of ITSM than can act to support 'management interoperability' in federated environments where disparate or competing organisations must cooperate to manage services.

A significant amount of effort was devoted to clarify the definition of services that are provided through EGI Production infrastructure. This involved the creation of a new dedicated board, the Services and Solutions Board.

In addition to this, EGI revised its engagement process: after the pre-production phase, a Service Level Agreement is established for each new user community, such that the research community participates in the process of estimating its service requirements, and service providers (EGI

participants at national/international level and Resource Centres locally) are engaged in committing to service levels and long term service provisioning. The new process is already being applied in the engagement with user groups. EGI is working with a number of organizations, research communities and VRE operators (DRIHM, BILS[27], Terradue, Mobrain[28], Pancancer, Life Science Grid Community, iMARINE, EXTRAS project, Human Brain Project and the nanotechnology research community) to establish a Service Level Agreement (SLA) with resource providers. SLAs are not legal contracts but, as agreements, they outline the clear intentions to collaborate and support research. To support this work a process of SLA negotiation has been defined and followed to ensure effectiveness and repeatability of this activity.

Once an SLA is agreed, EGI continues to support the effort between the resource providers to enable the research community on the promised resources as well as future monitoring.

The base for SLA negotiations with research communities are the Resource Centre (RC)[29] and Resource infrastructure Provider (RP)[30] Operational Level Agreements established in 2014. Both documents ensure that resource providers are properly integrated with the EGI Infrastructure and provide minimum required availability and reliability of resources.

[27] EGI BILS VO SLA and OLAs https://documents.egi.eu/public/ShowDocument?docid=2701

[28] EGI MOBRAIN SLA and OLAs https://documents.egi.eu/document/2751

[29] https://documents.egi.eu/public/ShowDocument?docid=31

[30] https://documents.egi.eu/public/ShowDocument?docid=463

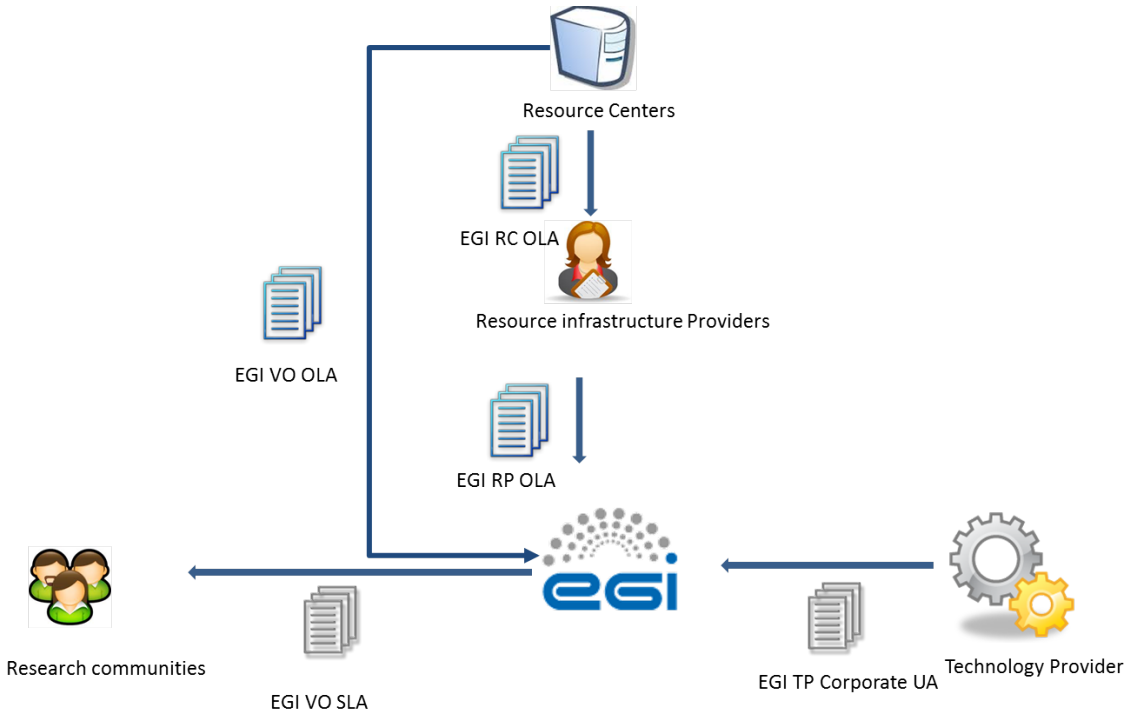

Figure 39 EGI OLA and SLA scheme

The figure above illustrates the process to negotiate the SLAs/OLAs that engages providers, EGI.eu and the research communities. EGI negotiates centrally the SLA with the research communities, liaising with the NGIs/EIROs and the resources centres. The outputs of this internal negotiation are the VO-specific OLAs who are extending the RC OLA and RP OLA; these documents support the VO SLA, ensuring that the targets of this last agreement are fulfilled by the service providers.

Another area the project has been improving is suppliers and customer relationship management where suppliers and customers for each service have been identified and unified approach has been defined to manage performance, satisfaction and complains.

To support production tools release and deployment process a new procedure is in preparation which aims to define steps that needs to be performed before new release of services under Software and Service Platforms category can be introduced into production infrastructure. Thanks to this procedure it will be ensured that each release is properly tested and documented, and will not impact the Infrastructure in a negative way.

In the coming year EGI will work towards implementation of the remaining processes defined in the FitSM standard:

- Service Reporting Management
  - Defining all service reports and ensure they are produced according to specifications in a timely manner to support decision-making.

- Service Availability & Continuity Management
  - Ensuring sufficient service availability to meet agreed requirements and adequate service continuity in case of exceptional situations
- Capacity Management
  - Ensuring sufficient capacities are provided to meet agreed service capacity and performance requirements.
- Information Security Management
  - Managing information security effectively through all activities performed to deliver and manage services, so that the confidentiality, integrity and accessibility of relevant information assets are preserved
- Problem Management
  - Defining way to investigate the root causes of (recurring) incidents in order to avoid future recurrence of incidents by resolving the underlying problem, or to ensure workarounds / temporary fixes are available
- Configuration Management
  - Defining way to provide and maintain information about all services and their relationships and dependencies
- Change Management
  - Ensuring changes to the services are planned, approved, implemented and reviewed in a controlled manner to avoid adverse impact of changes to services or the customers receiving services

4 Evolution of the security operations

Security operations, policies, procedures and best practices, have been evolved during PY1 to meet the requirements of new trust models, new developments and new usage scenarios of EGI services. As usual, developments in policies and procedures have been driven by risk assessment of the security requirements and trust models of the new EGI services, including the use of Federated Identity Management, the EGI Federated Cloud platform and the services for the Long Tail of Science.

In this section we present the results achieved so far and the planned ones to take place in PY2.

4.1 Security policies

A meeting of the EGI Security Policy Group at the start of EGI-Engage considered which security policies were most in need of revision to address the new EGI-Engage services. It was decided that the "Grid Acceptable Use Policy" and "The Security Policy for the Endorsement and Operation of Virtual Machine Images"[31] should be addressed by EGI-Engage in its first year. This work has been done in parallel with the production of new policies for the Long Tail of Science service (LToS) allowing for a simplified identity vetting procedure for users that do not belong to a community VO, and of a draft of a new general policy addressing Data Protection issues. All of this policy work is required to ensure managerial controls on the operation of new services developed in EGI-Engage, thereby mitigating the related security risks.

The new AUP, now called "Acceptable Use Policy and Conditions of Use"[32], has been generalised to include all EGI service offerings (Grids, Clouds, LToS, etc.). At the same time wording was changed to require appropriate acknowledgement of use of resources and support in publications. The policy on VM Endorsement[33] has been modified to better fit the policy and trust issues in the EGI Federated Cloud service.

A large number of users and communities would benefit from expanded use of credentials based on Federated Authentication ("FedAuth") for accessing services instead of employing personal certificates either installed in browsers or used to create proxy certificates. Secure use of FedAuth on the infrastructure relies on two capabilities:

- a **translation of FedAuth credentials** to a form understood by the services
- some agreed mechanism to retain the current credential assurance levels and confidence in identity vetting.

[31] https://documents.egi.eu/document/771

[32] https://documents.egi.eu/document/2623

[33] https://documents.egi.eu/document/2729

In this context, it is important to realize that user communities may be structured in different ways, and that, depending on the internal structure, coherence, and level of organization of the user community, it may or may not be able to provide user traceability, assurance, or identity vetting. Of the many user communities supported by EGI, only a small sub-set is so stringently organized as to be able to independently perform high-assurance identity vetting. In other cases (in particular in the context of the specific provisions for the Long Tail of Science), EGI centrally takes a responsibility of providing higher-assurance identity and traceability information independent of the user community. In the majority of cases, the communities actually rely on external identity vetting information in enrolling their members.

Only for highly organized communities, and for those cases where EGI independently provides assurance, FedAuth can be leveraged early since only minimal information (a persistent non-reassigned identifier) is needed from the identity providers in the FedAuth infrastructure. For those cases, and for the first concrete implementation of a prototype around the 'Federated Identity Management for Research' (FIM4R) pilots (e.g. the WLCG WebFTS use case), a co-existence model and policy needs to be in place in order to support both these highly structured as well as the more dynamic communities on the same infrastructure. To advance the development of a coordinated trust policy that will in the future be able to more dynamically accommodate differentiated trust models, WLCG in collaboration with EGI has developed a co-existence model and evolution for a sustained implementation. This model was developed with support from EGI-Engage and is documented in the evolving document "Considerations on the coexistence of controlled and flexible community models"[34].

Also the "Long Tail of Science" specific policy leverages differentiated and redistributed responsibilities. Here the policy distributed specific elements of the end-to-end risk assessment to the registrars within the EGI community, the centralized "User Management Portal" of the EGI LToS service, and the participating resource centres, aiming to contain any residual risks exposed through the LToS service towards other, non-participating resource centres and NGIs. The policy aims to enable a low-barrier Service to be offered to a wide range of research users in Europe and their collaborators world-wide, by any Resource Centre organisation that elects to do so. In offering such LToS Services, the Resource Centre shall not negatively affect the security or change the security risk of any other Resource Centre or any other part of the e-Infrastructure. In particular, security incidents originating in the LToS Service should not impact the IT Infrastructure in ways that are incompatible with the operational model of other, more tightly controlled, parts of the infrastructure. This document also provides guidelines on the implementation of security procedures and controls to facilitate offering of the Service by Resource Centres and Science Gateways.

The Guidelines[35] also contain normative information on how to implement the Policy.

[34] https://documents.egi.eu/document/2745

[35] https://documents.egi.eu/document/2734

A version of the new AUP specific to LToS[36] has also been produced and adopted.

Future work on security policies during year 2 of the project will be aimed at making other additional policies more applicable to the new EGI-Engage use cases. This will include a revision of the top-level overall Security Policy document and a revision of the Virtual Organisation Membership Management Policy to apply to the wider range of user communities now being addressed in EGI-Engage.

The EGI policies will also continue to be updated to reflect the new European regulations that apply to the EGI services. European directives and regulations have been always used as important input in the definition of the EGI security policies, and in the future this will be more important as EU increasingly focuses on the secure provision of IT services part of the EGI portfolio, such as cloud[37]. The deployment of new data services will require to further review the data protection policy in order to align the EGI policy to the EU directives and to provide trustful services to the users.

4.2 Security procedures

In order to provide efficient Operational Security in evolving infrastructures the security procedures constantly have to be developed further.

The technological aspect of the procedure development aims to exploit new possibilities in incident response required by the newly integrated technologies. In addition the new players (e.g. cloud resource providers) have to be integrated into the overall incident response concept. This activity is reflected in the EGI CSIRT Security Incident Handling Procedure[38]. This procedure was presented to OMB for approval.

To maintain a properly patched infrastructure and make sure that CRITICAL Vulnerabilities are handled adequately by all involved entities the EGI-CSIRT Critical Vulnerability Handling procedure[39] was further developed. This procedure is currently in draft, here the new developed supporting policies need to be approved.

In addition to the procedure development, also security monitoring and incident response tool development will be addressed in project year 2. The aim is to be able to monitor and enforce the policies developed here and to preserve the central user management capabilities as for the extended set of services provided by EGI. The EGI Software Vulnerability group vulnerability issue handling procedure has been revised and approved by the EGI Operations Management board.

[36] https://documents.egi.eu/document/2635

[37] One example is the Network and Information Security (NIS) directive: http://www.consilium.europa.eu/en/press/press-releases/2015/12/18-cybersecurity-agreement/

[38] https://wiki.egi.eu/wiki/SEC01

[39] https://wiki.egi.eu/wiki/SEC03

Previously during EGI-InSPIRE the main focus of the EGI issue handling was on the Grid Middleware distributed in the EGI UMD, and additionally to assist EGI CSIRT in the risk assessment of other software vulnerabilities, mainly in the Linux operating system. Technology is changing, in particular related to the emergence of the EGI Federated cloud. A much wider variety of software is in use such as Cloud enabling software, software within VMs, VMs themselves, VO specific software. SVG cannot control what software is in use. Some of this software is commercial; produced by large or small companies or organisations. Some is produced by EGI partners. Some software is released in the EGI UMD by resource providers with which EGI has a service level agreement; some such as operational tools for EGI infrastructure is released by the EGI team, as well as VOs, which take their software from a much wider variety of sources. This means we needed to revise the way we minimize risk arising from software vulnerabilities to the EGI infrastructure.

The advisory template has also been revised, taking account of comments from various site administrators, including so that the basic information and what is required of sites is displayed in the e-mail preview page and to make it more mobile friendly.

It has been found that various Virtual Organisations and user groups have been developing or using software which is not as secure as we would like, or configured in a way that is not secure or compliant with policy. Also this wider variety of software means the SVG cannot be experts on much of the software deployed in the EGI infrastructure. For this reason we included a Software security Checklist to help those who are developing or selecting software avoid some of the most common problems. This has been made available on the EGI Wiki[40]

For technology on which the EGI federated cloud heavily relies, a Technology Provider questionnaire was produced to ensure that the technology is reasonably secure and suitable for use in the EGI infrastructure. The idea is that this questionnaire is filled in for any technology which is deployed on the EGI infrastructure, and on which EGI relies. This provides some assurance that at least at the time this questionnaire is filled in, it does not contradict EGI security policy. It is not a full security analysis of the software. This questionnaire may be filled in by the developers, which is the case for software being developed for use on our infrastructure. Or it may be filled in by someone who is selecting a technology, or who has expertise in that technology in EGI.

A version of the Technology provider questionnaire was approved at the EGI Operations management board in September 2015.

4.3 Security risks assessment

EGI Security Threat Risk assessment with focus on the EGI Federated cloud and the changing EGI environment is being carried out at the time of writing and is near completion. A similar approach to that in 2012 is being carried out. First a team of people to carry out the work was established,

[40] https://wiki.egi.eu/wiki/SVG:Software_Security_Checklist

this included people from CSIRT, the Software Vulnerability Group, the EGI Federated cloud, Security Policy Group, and others. Then a draft set of threats was produced, starting from the list from 2012 but adding new ones both from general experience since then and those associated with Virtualization and the Federated Cloud. These threats were divided into various categories, and members of the team were asked to take one or two categories each and improve on them, and provide a description of or update the current situation. Then the team were invited to comment, add others. A list of 103 threats was produced, in 18 different categories. Then all members of the team were invited to provide their opinion on the 'likelihood' and 'Impact' of each threat, according to certain guidelines. In all 10 members of the team returned a spreadsheet with their opinion of likelihood and impact, and the average risk was computed from this. At present, the report is being prepared, including suggestions for mitigation of some of the highest risk threats, and is due in 2016 Q2.

5 Roadmap for the EGI production infrastructure

The roadmap for the EGI Operations is developing in two directions:

- Consolidating the current production services, increasing their performances, and their reliability, with the final goal of improving the user experience.
- Integrating in the production infrastructure the new platforms and services that are produced by the EGI-Engage project and other activities and collaborations of EGI.

These two main themes are described in detail in the next sub-sections.

5.1 Consolidating current production services

The production services of EGI must continue to evolve with the goal to improve the user experience. All services are periodically assessed to understand where there are issues degrading the quality of service perceived by our users. The HTC services, considering the longer experience, have reached a good maturity level, with a considerably decreasing number of issues also helped by the fact that the middleware software is not being updated frequently, with major changes.

The federated cloud services are in production for less than two years at the moment of writing, and although in constant improvement the quality of the service provided to the user is not very uniform being dependent on local configurations and settings, and then can thus vary much depending on the service provider used at the moment. The following activities will be undertaken in PY2:

- Further improvement of documentation for the service providers and development of different configuration options in integrating via OCCI to reduce barriers of adoption.
- Development of new testing probes for monitoring of all cloud capabilities.
- Development of automated testing of the components integrating different cloud management systems in EGI.
- Application of the EGI quality verification and release processes to the “plugins” developed by EGI for the cloud management system aiming at enhancing the stability of new releases and at detecting malfunctions related to the diverse deployment environments of the EGI federated cloud.

5.1.1 Documentation

At the moment the documentation is fairly complete, what is partially missing is a top-down structure that would lead the cloud sites from the first approach to EGI through the integration and certification process. In particular documentation should highlight the different deployment scenarios in order to help the cloud provider in choosing the best architecture to federate with EGI. A revision of the current documentation structure is in progress at the moment of writing.

5.1.2 Monitoring

The automatic monitoring of the services is covering the basic functionalities of the services, for example virtual machine instantiation and removal. The user experience can though be affected by a number of other factors. One example is the availability of the virtual machine images, or advanced contextualization features that are needed by the use cases, whose distribution at the moment are not monitored.

The current plan is to start a detailed assessment of the status of the sites in the federated, through manual testing actually replicating the work of the users, at least touching most of the functionalities expected by users. This will achieve two results: sites with issues will be asked to fix their issues, and recurring issues will generate monitoring probes that are closer to user intended behaviour.

5.1.3 Integration testing with the cloud management system and software packaging

At the moment, to federate private clouds, EGI is developing and maintaining a set of extensions to the community based cloud management system (e.g. Open Stack and open Nebula). Maintaining these components requires keeping the compatibility with the new version of the cloud management system, and makes the EGI development available to the site managers who need to install them.

The integration tests with new CMS are currently implemented with some of the components, for example the OCCI integration with OpenStack is tested by the developers as soon as a new version of OpenStack is pre-released by the OS collaboration, but other components are not validated with the new releases of the main CMS, not with a reliable and repeatable process at least. This will have to improve with dedicated testbed resources, and automated testing procedures to validate new developments on both sides (EGI and the CMS).

Good testing of new releases will be the main building block for a proper distribution of the releases to the EGI sites. The distribution in UMD, or equivalent process, of the federated cloud components will force the developers to achieve: proper packaging and testing, clear documentation, plus a trusted single entry point for the download of the federating software.

Besides the federated cloud, another improvement deployed in the coming months is a centralized monitoring, that will free the NGIs from the burden of deploying a dedicated Nagios instances to submit probes to their sites, by deploying a centralized group of Nagios services that can support the whole infrastructure. This will reduce the maintenance cost of the Nagios services, and at the same time it will make much more flexible the monitoring system, which will not require a deployment campaign of a new version to add new probes or similar changes that now require NGIs intervention.

5.2 Integrating new services in production

In order to integrate new services in EGI, some prerequisites must be fulfilled:

- Technical support for the EGI AAI infrastructure
- Support for the EGI Accounting infrastructure, if relevant
- Monitoring probes to be integrated with the EGI monitoring infrastructure
- Fulfilment of EGI security policies

In some cases, it is acceptable to have a roadmap for the development of the requirements above, if the services need to be quickly integrated and made available to the users. Service provider must sign an OLA with EGI, where the targets for the service are defined.

Besides technical integration, EGI must ensure that the quality of the software supporting the services is production-ready. This is implemented by ensuring the software support by the developers, through an underpinning agreement, as described in section 3.4 and through the software quality assurance process described in section 3.3.

In EGI-Engage a number of services are being developed. Many Competence centres (CC) will integrate high-level discipline specific services targeting the represented communities. These services, or part of them, will be technically integrated and offered as EGI services to a wider community. This is one of the changes foreseen for the EGI production infrastructure, the users will not only access the resources, through services providing plain computing or storage, but will use also platforms that, underpinned by EGI resources, will provide high level services.

The new platforms can be of different level of abstraction, from a PaaS to a Virtual Research Environment, operated by EGI partners, or by the VOs.

To support these new platforms, EGI Operations must focus the effort in different directions, for example to develop the security policies and processes to assess the policy compliancy of the new technologies and access paradigms, and to properly react to any security issue that could involve the new platform, as it is done for every EGI service.

At the moment the operational infrastructure is hierarchically organized in NGIs/EIROS and resource centres. Every service is provided by a resource centre, and every resource centre is connected to an NGI. If EGI integrates services offered by other entities, e.g. communities or external providers, the hierarchical structure will have to be extended to include in the operational framework other entities providing services part of the EGI portfolio.

While the technical implementation of monitoring should not change from any development done for the current services, the reporting, how the monitoring results are used, will have to change. Currently any deviation in performance is handled by the NGI and by the site staff, but adding on top of these services a platform operated by – let's say – a VO, introduces a new level of complexity. Availability of the platform can be monitored and reported, to the entity operating the

service, as well as the supporting resources, can be monitored and reported to the resource centres operating the resources.

The development of the EGI Marketplace will push the boundaries of the scenario described above. On top of what the CCs will integrate in EGI, through the marketplace a number of new services will be added. While at the moment of writing the structure and the policies behind the marketplace are not yet defined, we can anticipate that the services accessible through the Marketplace will be very diverse, for example in the form of virtual appliances and datasets. As every service provided by EGI, Operations will have to work with the providers in order to understand how service availability can be monitored, reported, and how targets should be set, and ultimately upon which events EGI Operations can remove a service from production, preventing users to access it. Moreover, at least a subset of the services in the marketplace will have to be monitored and the status information be made available to the users who want to access them.

The evolution of the AAI infrastructure, under the JRA1 work package, is one of the topics of EGI-Engage with the biggest impact on operations at all levels. The most important impact is the new capabilities offered to the users, who will be able to use their institutional credentials to access EGI, and this will affect all the layers of the production infrastructure. Once the AAI layer will be enabled for federated authentication the current EGI services must implement the support for the new authentication services.

Ultimately, AAI layer is composed by a number of components that need to be rolled in production, and integrated in the set of core services provided by EGI to our stakeholders. The deployment of these production instances must be planned, and appropriate resources be allocated, but also the components must be monitored and supported by appropriate OLAs and policies to be production-ready.

The other big change that is being introduced by EGI-Engage is the Open Data Platform. Similarly to other new services, the outputs of the JRA2 work package will have to be integrated in the operations framework. But beside the operational details, the new set of services of the open data platform will be a big shift in the EGI service provisioning, adding to the current one, compute-centric, a new data-centric resource provisioning. Users will be able to use EGI as a data infrastructure, and use computing associated to data. These features will be supported by appropriate monitoring and accounting, and offered through the marketplace.

Some of the current and the new services will be also offered through the long tail of science, which is an access mode to get basic services, and limited resources, without the overhead of setting up a virtual organization, or ask for a grant of resources. The long tail of science access mode can be also considered a tool to offer a “try before you buy” mode, that could consequently evolve in a long-term collaboration with a community.

Ultimately EGI production infrastructure will hopefully incorporate also the outputs of other activities with a strong collaboration with EGI, such as the INDIGO DataCloud project, which will

release both infrastructure/service providers' oriented developments and PaaS and SaaS for the user communities.