EGI-InSPIRE

PERIODIC REPORT

Document identifier:	EGI-InSPIRE-PR-PY2.odt
Date:	**06/07/2012**
Activity:	**NA1**
Lead Partner:	**EGI.eu**
Document Status:	**DRAFT**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/1179

Abstract

The following report provides the periodic report of the EGI-InSPIRE project at the end of its second year. The report provides details of the scientific achievements that have been obtained over the period which provides an overview of the details contained in the three quarterly reports issued at the end of PQ5, PQ6 and PQ7.

I.COPYRIGHT NOTICE

II.DELIVERY SLIP

	Name	Partner/Activity	Date
From	Steven Newhouse	EGI.eu/NA1	1/6/2012
Reviewed by	**Reviewers:**	PMB	
Approved by	**AMB & PMB**		

III.DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	10/6/12	First draft	Steven Newhouse/EGI.eu
2			
3			

IV.APPLICATION AREA

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

V.DOCUMENT AMENDMENT PROCEDURE

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

VI.TERMINOLOGY

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

VII.PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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

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

TABLE OF CONTENTS

Publishable Summary 7

Project Progress 12

Project Objectives for the Period 12

Work progress and achievements during the period 14

Project Issues 88

Project Management 89

Deliverables and Milestones Tables 93

Explanation of the Use of Resources 98

Summary 98

Resources consumed per beneficiary 101

Financial Statements per Beneficiary 102

Summary 102

Certificates 103

Annex 1: Dissemination and Use 104

Main Project and Activity Meetings 104

Conferences/Workshops Organised 105

Other Conferences/Workshops Attended 110

Publications 121

1 Publishable Summary

During its second project year EGI-InSPIRE continued with the operation of the a European wide distributed infrastructure encompassing over 350 resource centres across over 50 countries that supports over 21,000 multi-disciplinary researchers in their everyday data analysis work. Within the project's activities, as a result of the recommendations coming from the first EC review, there was a significant reorganisation and a renewed focus on the work being undertaken on non-operational activities supported by EGI-InSPIRE. At the heart of this reorganisation was the establishment of a flexible dynamic mechanism to harness expertise within the NGIs and EGI.eu to tackle issues of community wide importance. Through the establishment of dedicated national contact points – the NGI International Liaison (NILs) role – a human network was established that could bring experts within an NGI into Virtual Teams. These virtual teams are created for 3-6 months and on average encompass representatives from under a dozen representatives from different organisations. By the end PY2 two VT projects had been completed, nine were in progress and a further seven projects are in the process of being established.

Operational Infrastructure

The transition started in PY1, from regional federated to national independent Operations Centres was completed. The total number of Resource Centres (RCs) in March 2011 amounted to 352 instances organised into 27 national operations centres and 9 federated operations centres. Availability and Reliability reached 94.50% and 95.42% (yearly average), which amount to a +1.8% increase of availability and a +1.5% increase of reliability in PY2. Overall resource utilization has been satisfactorily progressing confirming the trends of PY1. The total number of jobs executed in EGI between May 2011-April 2012 increased by 46.42% of the yearly job workload done from May 2010 to April 2011.

The EGI.eu central tools were significantly advanced during PY2. The Metrics Portal was placed into production. The message broker network was repeatedly upgraded to improve the reliability of message delivery, stability, manageability and scalability. The transition from R-GMA to messaging of the accounting infrastructure was completed and a new central consumer based on ActiveMQ STOMP was deployed in pre-production. New accounting portal releases provided role based views and new graphics. GOCDB functionality was also significantly extended with the support virtual sites, new roles and permissions, scoping of RCs and service end-points, and a hardened DNS-based failover configuration. The Service Availability Monitoring (SAM) underwent seven different upgrades and is currently the largest and more distributed operational infrastructure comprising 32 distributed instances. The operations portal rolled to production new major components: the VO Dashboard and the Security Dashboard. In addition the VO management features were greatly enhanced. The EGI Helpdesk's support units were updated, a new report generator was designed and prototyped, FAQs were migrated to the EGI wiki, usability of the system was enhanced is now interfaced to a new helpdesk system (Service NOW). The failover configuration was hardened with auto-switching between different front-ends. A new set of catch-all services for the monitoring of uncertified Resource Centres was placed into production.

EGI's Central Grid Oversight team in addition to its certification and management duties was involved in training and dissemination activities, in follow-up of underperformance both at a RC and at a Resource Provider level, and in monitoring the instability of the distributed SAM infrastructure. EGI's operational security team ensured day-by-day security monitoring, and timely response in case of incidents. The EGI Security Threat Risk assessment identified 75 threats in 20 categories and an initial

risk assessment and preliminary report was produced describing the assessment process, progress and initial findings. Specialized tools for incident response tracking and for streamlining of operational security tasks were prototyped and rolled to production. The Staged Rollout workflow introduced in PY1 was progressively refined and gradually expanded reflecting the increasing deployment needs of the community, and resources were reallocated to ensure testing of a broader range of products. Participation of NGIs in Staged Rollout activities as Early Adopters has been increasing, but distribution of load across the NGIs is now being revised, as with the release of EMI 2 different platforms are being supported (SL5, SL6, Debian) and Staged Rollout effort needs to be concentrated on products whose deployment is requested at higher priority.

The operations integration of GLOBUS, UNICORE, QosCosGrid and Desktop Grids was completed, with the exception of accounting, which requires further integration development. Extensions are being implemented in collaboration with the relevant external technology providers. Documentation was revised and migrated to the EGI wiki, three new operational procedures were approved and training and support pages were improved.

The EGI Operations Level Agreement framework was considerably extended in PY2 with the first Resource Centre Operational Level Agreement, defining the target levels of the services provided by sites for resource access, and the Resource infrastructure Provider Operational Level Agreement, defining the target levels of the community services provided by the NGIs, which came into force in January 2012. Several NGIs and RCs have been frequently experiencing low performance in PY2. Problems differ considerably, among these the unreliable infrastructure of the machine room hosting resources and services, unavailability of expertise in running grid middleware, lack of experience with procedures and best practices. A support plan for some of the NGIs belonging to the South-East European region was approved by the project. GRNET is now supporting Armenia, Montenegro and FYRoM to overcome some identified technical issues, and training activities are being organized in collaboration with the technology providers.

Continued maintenance and third level support of software in EGI is paramount. Both will be challenged by the upcoming end of the EMI and IGE projects with discontinuation of maintenance and support for as yet undefined subset of products, lower quality of the support that is currently subject to SLAs, phasing out of the external repositories, and a change in the software distribution processes that will require changes in the EGI software provisioning processes.

The major goal for the software provisioning activity in PY2 was to put the software provisioning infrastructure into action that was designed and tested during PY1. When EMI-1 was released, more effort was required in May and June 2011 to make the processes work than had been initially expected. Releases are assessed by the Quality Criteria management team which has now defined more than 200 EGI Quality Criteria which are regularly reviewed by Technology Providers, Platform Operators and Integrators, and User Communities. Seven Unified Middleware Distribution releases were published during PY2 which comprised more than 100 initial product releases and updates for 56 products. The established processes of the Deployed Middleware Support Unit have proven to work well as the experience and expertise of the staff improves. The staff will now manage and monitor estimated arrival dates of key software updates coming from Technology Providers that will be used to monitor the accuracy of the estimates and through this the experienced maturity and reliability of change management processes.

Heavy User Community Support

During PY2 support was provided to EGI's "Heavy User Communities" through domain-specific support for the High Energy Physics, Life Sciences, Earth Sciences and Astronomy and Astrophysics communities together with support for Shared Services and Tools that cover not only these but also other user communities. A closer working relationship has manifested itself through the EGI Forums with much more pro-active interaction between the different communities as well as focused workshops covering sustainability models and common requirements in the area of long-term data preservation coupled with the continued ability to use the data. The long-established model of collaborative support, whereby responsibilities are spread over multiple institutes and organisations, continued to be the most appropriate sustainability model as it provided a satisfactory but not complete degree of isolation from shorter-term staffing and budget issues at the individual members of any given collaboration. One of the key successes of the work package – in terms of increased commonality of services and tools – is expected to lower the long-term support needs from what they would otherwise have been.

Community Engagement

A new website (www.egi.eu) featuring new images, a new structure and enhanced social media interactivity was launched in March 2012 to be the centre point of EGI's community engagement activities. A number of use cases have been developed and made available on the website and published through other channels, such as blogs and articles. From the 1 May 2011 to 30 April 2012, over 82,000 people visited the EGI web site, corresponding to 199,400 visits and 814,000 page views. Since the start of the project, there have been over 100,000 unique visitors to the website, with more than a million page views.

Internal and external communications activities include four issues of the EGI Inspired newsletter, monthly Director's letters and published articles about EGI in the e-IRG Newsletter, *The Parliament* magazine, the CLARIN newsletter, *Public Service Review: European Science & Technology*, *Supercomputing Online*, *HPC in the Cloud*, *International Innovation*, *Public Service Review: European Union, PanEuropeanNetworks: Science & Technology* and 15 EGI-related items in iSGTW. Some items have been picked up by the general media, such as *Symmetry*, *Discover News* and *Wired*.

Dedicated social media feeds for general, user community, tech/ops and policy have been established, which have been used to launch community competitions. EGI has produced a series of videos called "Stories from the grid". The first video "Episode 1: The cone snail and the hunt for more power anaesthetics[1] " was released on 29 February 2012.

EGI has hosted booths at or attended many events during the year, some of which have attracted up to 10,000 delegates, including the HealthGrid 2011 event in Bristol, eChallenges in Florence in October 2011, SuperComputing'11, Seattle, US in November, CloudScape IV in Brussels in February, ISGC2012 in Taiwan, ICRI2012 in Copenhagen and the EGU General Assembly in Vienna in April, producing dedicated materials for most events.

The *EGI Technical Forum 2011* was held from 19-23 September 2011 in Lyon, France. The event was organised in collaboration with local hosts CC-IN2P3 and France Grilles. In total, 655 participants registered for the event. There were 132 contributions from 296 speakers and 34 session conveners. The event was collocated with the Open Grid Forum, Grid2011, GlobusEUROPE, a French Grid Day and the 9th e-Infrastructure Concertation meeting, which was organised by the e-ScienceTalk project in collaboration with the EC. The event was attended by a journalist from *HPCinthecloud* and *Datanami*,

[1] http://go.egi.eu/conco

leading to several articles. During the event, there were 250 Tweets from 27 people, 20 photos on Flickr, 27 blog posts on GridCast, including 9 videos. A feedback survey was launched at the end of the event and received 114 responses. Nearly 90% of respondents found the conference website quite or very useful. Around 75% found the registration process quite or very easy to use and around the same percentage found the EGI staff helpful. For the social media channels, around a quarter referred to the EGI blog, 14% to iSGTW, 43% read the GridCast blog, 19% looked at the Facebook group and 45% used Twitter.

The *EGI Community Forum 2012*, was held in Munich on 26-30th March 2012, in conjunction with the 2nd EMI Technical Conference and co-located with the 2nd Annual European Globus Community Forum. The event was advertised in *HPCwire*, *HPC in the cloud* and *Datanami* and journalists attended from iSGTW and the US. The event was attended by 518 participants and featured 171 contributions from 208 speakers. During the event, there were over 400 Tweets from more than 60 users. Around 220 photos were uploaded to Flickr and over 150 delegates downloaded the Conference4me conference app. The GridCast blog brought together 10 bloggers, posting more than 30 posts and 14 videos from the event. The event website received over 5380 visits from 1700 visitors. From 122 survey responses, 93% found the conference website useful or very useful and 96% reported that registration was easy to use. 90% reported that the conference staff were helpful, or very helpful. For social media channels, of those that responded, 52% used Twitter, 25% Flickr, 18% viewed YouTube, 30% the GridCast blog and 45% the EGI blog.

The Strategic Planning and Policy Support activity helped shape EGI's 2020 strategy and sustainability plans. Work included the analysis and definition of the EGI ecosystem and defining a business modelling framework for the key roles participating within the ecosystem. EGI's potential contributions to the European Commission's EU2020 policy has been analysed and key actions have been identified and been used as the basis for input to a number of survey and position papers have been provided to influence policy makers. Strategy and policy messages have been regularly delivered through the different media channels and key workshops have been organised at the EGI events. Important security policies and procedures have been developed and approved over the year and 16 new MoUs have been signed to consolidate the external strategic partnerships needed by EGI. All the agreements have been continuously monitored to track the proper execution. Two new initiatives have been successfully executed: a common EGI glossary to harmonise the usage of key terms across the various functional areas of EGI and the establishment of the EGI Compendium to set up a process and structured information to be collected from NGIs/EIROs and develop a strategic body of information for the EGI management bodies and for the community at large.

The Community Outreach team is focussed on establishing and maintaining communication with structured user communities in order to identify and address their distributed computing-related needs by attending user community meetings and organising workshops at the EGI Forums. Three new Virtual Research Communities (VRCs) were recognised through signed Memorandums of Understanding with the: Life Sciences Grid Community; Hydro-Meteorology Research Community, Worldwide LHC Computing Grid and the Humanities represented by the CLARIN and DARIAH ESFRI projects in the form of a Letter of Intent. The User Community Board meetings provided a regular forum for user community representatives to engage with EGI and its various planning activities. Considerable effort went into engaging with the ESFRI projects and others to enable them to participate and share their technical needs with the EGI community. As a result of these discussions various partnerships and working relationships have been established.

The Technical Outreach New Communities team (established as a separate focused activity in PM19) has worked in close collaboration with other EGI.eu teams and NGI staff in a number of areas. The development of EGI's Training Marketplace, Application Database, Virtual Organisation Services and the establishment of a Contact Relationship Management System to facilitate systematic approaches to new user communities continued. In addition, workshops developing and supporting the adoption of gateway and other services have been run, technical support pages for users and developers created, content within the Application Database updated and the user requirements analysed and processed. The team has used the blog to propagate support-related information and played a leading role in establishing the Virtual Team model through driving forward the overall coordination as well as running a number of VT projects.

2 Project Progress

2.1 Project Objectives for the Period

EGI-InSPIRE defines the following project objectives (PO) as its goals:

- **PO1:** The continued operation and expansion of today's production infrastructure by transitioning to a governance model and operational infrastructure that can be increasingly sustained outside of specific project funding.

- **PO2:** The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.

- **PO3:** The support for current heavy users of the infrastructure in Earth Science, Astronomy & Astrophysics, Fusion, Computational Chemistry and Materials Science Technology, Life Sciences and High Energy Physics as they move to sustainable support models for their own communities.

- **PO4:** Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.

- **PO5:** Mechanisms to integrate existing infrastructure providers in Europe and around the world into the production infrastructure so as to provide transparent access to all authorised users.

- **PO6:** Establish processes and procedures to allow the integration of new DCI technologies (e.g. clouds, volunteer desktop grids, etc.) and heterogeneous resources (e.g. HTC and HPC) into a seamless production

Performance of the individual activities through their activity metrics (defined in D1.9[2]) can be found in the metrics portal[3]. The performance of the project against the project metrics (defined in D1.7[4]) are summarised below.

[2] https://documents.egi.eu/document/1071

[3] http://metrics.egi.eu

[4] https://documents.egi.eu/document/1019

Table 1: Achieved PY2 Project Metrics (PQ5-PQ8)

Project Objectives	Objective Summary	Metrics	PQ5	PQ6	PQ7	PQ8	Target PY2
PO1	Expansion of a nationally based production infrastructure	Total number of production resource centres in EGI (M.SA1.Size.1)	346	348	347	347	330
		Total number of job slots available in EGI-Integrated and peer infrastructures (M.SA1.Size.2)	337,608	364,500	399,300	419,274	350,000
		EGI monthly reliability of site middleware services (M.SA1.Operation.5)	95.3%	94.5%	94.8%	94.4%	91%
PO2	Support of European researchers and international collaborators through VRCs	MoUs with VRCs (M.NA2.11)	1	1	1	0	10
		(Cumulative) Number of papers from EGI Users (M.NA2.5)	39	58	85	85	60
		Average number of jobs done per day for all VOs (M.SA1.Usage.1)	1.029.830	1.331.525	1.264.922	1,350,000	525.000
PO3	Sustainable support for Heavy User Communities	Number of sites supporting MPI (M.SA1.Integration.2)	93	91	108	109	100
		Number of users from HUC VOs (M.NA3.12)	9.672	9.861	10.856	10,966	5.500

Project Object-ives	Objective Summary	Metrics	PQ5	PQ6	PQ7	PQ8	Target PY2
PO4	Addition of new User Communities	Amount of integrated desktop resources (M.SA1.Integration.3)	1	0	0	0	5
		Number of users from non-HUC VOs (From M.NA3.12)	5.840	6.413	8.518	7,269	1.000
		Public events organised (M.NA2.6)	6.491	4.870	11,975	14,047	2.000
PO5	Transparent integration of other infrastructures	MoUs with resource providers (M.NA2.10)	0	1	1	1	5
PO6	Integration of new technologies and resources	MoUs with Technology providers (M.NA2.9)	0	0	0	0	4
		Number of production HPC clusters (M.SA1.Integration.1)	56	38	39	32	3
		Amount of virtualised installed capacity accessible to EGI users (HEPSPEC) (M.SA1.Integration.4)	1	6	7	7	1

2.2 Work progress and achievements during the period

2.2.1 Operations

During PY2 SA1 was responsible for the continued operation and expansion of the production infrastructure. The transition started in PY1, which evolved the EGEE federated Operations Centre into independent NGIs, was completed. The total number of Resource Centres (RCs) in March 2011 was 352 instances (+3.22% yearly increase). The installed capacity and Resource Centres grew considerably to comprise 270,800 logical cores (+30.7% yearly increase), 2.96 Million HEP-SPEC 06 (+49.5%), 139 PB of disk space (+31.4%) and 134.3 PB of tape (+50%).

EGI currently contains 27 national operations centres and 9 federated operations centres encompassing multiple NGIs. Availability and Reliability reached 94.50% and 95.42% (yearly average), which amounts to a +1% increase in PY2 over the year. Overall resource utilization has been satisfactorily

progressing confirming the trends of PY1. The yearly increase of the total number of jobs executed in the infrastructure in the period May 2011-April 2012 amounts to a 46.42% increase of the yearly job workload done from May 2010 to April 2011. The PY2 overall quantity of EGI computing resources used amounts to 10.5 Billion HEP-SPEC 06 Hours.

The Operations Management Board (OMB) continued to provide the technical management of the production infrastructure: 11 meetings were organized of which 2 were co-located with the main project events. Operations liaisons with Open Science Grid in the USA were strengthened. Fortnightly operations meetings were run regularly to coordinate the technical work. SA1 run 5 task forces and coordinated the TCB task force on accounting.

The OMB contributed to the EGI sustainability workshop (January 2012) and to middleware sustainability discussions by providing information on status, costs and sustainability plans of NGI operations services, and by collecting priorities of the currently deployed software. The operations service portfolio was defined and documented in the EGI Operations Architecture: Grid Service Management Best Practices (D4.3[5]). This is important for future sustainability planning and the estimation of the operations cost.

2.2.1.1.1 Security

Operational security was run effectively during PY2 and ensured day-by-day security monitoring, and timely response in case of incidents. Security in EGI was reviewed following the PY1 reviewers' suggestions, and documented in D4.4[6]: the scope and aims of EGI security were reviewed, including the assets that EGI security seeks to protect. The work of the various security groups in or associated with EGI was presented. Practices and standards for IT security and their usage and possible future usage were analysed and plans for a security risk assessment were defined.

The EGI Security Threat Risk assessment team was formed and identified 75 threats in 20 categories. An initial risk assessment was undertaken and a preliminary report[7] was produced which identified 13 high-risk threats with a risk of 8 or more and 4 threats with a very high impact.

Specialized tools for incident response tracking and for streamlining of operational security tasks, were prototyped during the year and in some cases have already been rolled to production. These were introduced to support the security drills framework and were extensively used during the Security Service Challenge 5 (SSC5) – a cross-NGI security challenge. Various training events were organized, which were well attended.

Statistics. During PY2, 10 potential software vulnerabilities were reported to SVG, and SVG issued seven advisories[8]. EGI CSIRT reported on 10 security incidents, and issued four[9] security alerts, of which one was critical and two were high risk.

[5] https://documents.egi.eu/document/763

[6] https://documents.egi.eu/document/863

[7] https://documents.egi.eu/document/969 (Restricted Access - Available to EC Reviewers)

[8] https://wiki.egi.eu/wiki/SVG:Advisories

Procedures. The EGI Security Incident Handling Procedure[10] and the EGI Software Vulnerability Issue Handling Procedure[11] were both updated, and a new procedure, the EGI CSIRT Critical Vulnerability Operational Procedure[12], approved. The Operational Security Procedures report, is the entry point to the EGI operational security procedure framework and gives an overview of operational security processes[13].

Vulnerability Assessment of Grid Middleware. Security Vulnerability Assessment is the pro-active examination of software to find vulnerabilities that may exist. A Security Assessment Plan[14] was defined at the beginning of PY2 to identify which software components within EMI would be assessed and when the assessments would take place. The document also presents the status of software packages previously assessed.

Support tools. EGI CSIRT monitors the production infrastructure and notifies RCs exposing security vulnerabilities in order to correct weaknesses before security incidents occur. Results detected by security monitoring tool are made available to the RCs and NGIs. The EGI Operations Portal has been extended with a Security Dashboard[15], which aggregates information from the CSIRT monitoring tools and presents them in a uniform manner to authorized users. The first prototype was made available during PQ7 and was rolled to production in PQ8.

The security probes were extended to improve accuracy of tests and to limit the number of false positives. The Nagios instance that is responsible of security monitoring was moved under the egi.eu domain (secmon.egi.eu). Additional improvements have been applied to the Pakiti service, which detects known vulnerable packages that are installed at the RCs, mainly to improve its scalability and robustness in handling data from the whole of the production infrastructure.

A ticketing system to support incident response – RT for Incident Response (RTIR) was put into production.

Security Service Challenge (SSC). Purpose of EGI SSCs is to investigate whether sufficient information is available to be able conduct an audit trace as part of an incident response, to ensure that procedures are enforced and that appropriate communications channels are available and used. A SSC comprises activities in three areas: communication, containment and forensics. Performance of Resource Centres in these three areas is recorded. The role of EGI and NGIs in the SSC is to provide

[9] https://wiki.egi.eu/wiki/EGI_CSIRT:Alerts

[10] https://documents.egi.eu/document/710

[11] https://documents.egi.eu/document/717

[12] https://documents.egi.eu/document/283

[13] https://documents.egi.eu/document/649

[14] https://documents.egi.eu/document/563

[15] https://operations-portal.egi.eu/csiDashboard

coordination and to be a single point of contact to entities like VOs and Certification Authorities. Coordination includes making sure that that obtained forensic results are processed and that resulting information is fed back to the affected sites.

The SSC5 run was multi-site incident simulation exercise. A multi-site grid incident affects multiple NGIs and involves the entire EGI ecosystem (Resource Centres, NGIs, Certification Authorities, User Communities, security teams etc.). SSC-5 was therefore created as a realistic scenario involving 40 resource centres across more than 20 countries, and a VO-specific job submission framework was adopted (PANDA). A VO job-submission framework adds extra complexity to containment of the simulated incident on a Resource Centre level. SSC-5 was successfully completed in June 2011. A SSC framework was developed to scale up to a larger number of sites. This framework allows for: job-submission (using different methods), storage operations, the definition of a set of tasks: communication, User/Process management with target times, the automated generation of reports and scoring schema, recording of history and monitoring of progress.

The SSC-5 showed that EGI-CSIRT in collaboration with the VO-CSIRT and the CA is able to handle a multi-site incident. For efficiency purposes it is crucial for EGI-CSIRT to have a close collaboration with the VO-CSIRT and the Certification Authorities, which was very good during this exercise.

Security training. Various training events for Resource Centre administrators were organized aiming at spreading knowledge about strategies for incident prevention (security monitoring tools and system hardening and patching), handling and containment (incident response procedures and mechanisms to control access to grid resources), computer forensics and post mortem analysis of the compromised systems. Two security training sessions were organized at the Technical Forum 2011, and during the EGI CSIRT face to face meeting in April. The training was very well received and new training sessions are being prepared for EGITF 2012.

Coordination

EGI CSIRT activities are coordinated through weekly operation meetings and monthly team meetings (of which two have been face to face). As of May 2011 also the Security Vulnerability Group (SVG) have been holding monthly meetings to streamline coordination.

2.2.1.1.2 Service Deployment

Staged rollout

During PY2 staged rollout activities were reviewed to ensure the early adoption of software releases from the EGI main external technology providers (EMI and IGE). Resources (mainly concentrated on staged rollout of gLite 3.2) were reallocated for testing of EMI and IGE software releases.

The Staged Rollout workflow introduced during PY1, was refined through the experience gathered with the testing of EMI 1, this has been done in parallel with the construction of the Staged Rollout infrastructure, which is being gradually being expanded to reflect the deployment needs of VRCs and NGIs[16]. A new provisioning procedure[17] is now being explored by having two new UMD repositories:

[16] https://documents.egi.eu/document/478

[17] https://wiki.egi.eu/wiki/UMD:Provisioning:Proposal_SR

"untested" will contain all packages from the moment they are released by the technology providers will be used to perform the Verification step of the SW provisioning, and "testing" will be used to perform the Staged Rollout step. Those repositories are fixed and publicly available, allowing any site not committed to as Early Adopter for testing new versions of the software. Furthermore, more than one version of any software package may be on any of those repositories. This change in the procedure reduces the time lag between the release by a technology provider and the availability in the UMD repositories as new versions will be immediately available in the UMD testing repository. The testing repository will be disabled by default, and all sites will be able to point to UMD repositories for all services (only the base and updates repositories will be enabled by default).

In preparation for the release of UMD 1.0 in PQ5, 81 tests in total were performed resulting in 2 products being rejected. This number reduced in the following quarters as subsequent UMD updates only included the subset of products that were being updated. A total of 122 components were tested in total during PY2, of which eight were rejected. 192 staged rollout tests were undertaken, and the number of Early Adopter teams has been progressively increasing to test a growing set of products from EMI, IGE and EGI-InSPIRE JRA1 (operational tools). It currently amount to 60 teams. IGE 1 product releases underwent software provisioning for UMD1.2, broadening the range of products released in UMD. Two EMI products, namely VOMS/Oracle and FTS, were not included into software provisioning because of the very limited deployment scope (VOMS/Oracle), and because specific to one user community, who is already in charge of testing and staged rollout (FTS).

ARC and UNICORE are part of UMD through EMI; extensive work was carried out in order to have fully operational interoperability of sites deploying those products.

Operations integration. The full integration of ARC was accomplished in PY2, while the integration of GLOBUS and UNICORE were almost completed by the end of PY2. Two task forces dedicated to GLOBUS[18] and UNICORE[19] involving interested NGIs and technology providers were constituted in PQ5. The GLOBUS and UNICORE integration are waiting for accounting integration to be completed. The UNICORE integration required extensions to accept all RFC 3986 characters in GOCDB, and both UNICORE and GLOBUS Nagios tests were integrated into the Service Availability Monitoring (SAM) Update 13.

GLOBUS accounting integration depends on the availability of GridSafe, whose extended version supporting messaging, will be shipped for the first time with the IGE 3.0 release expected in September 2012. UNICORE accounting integration is in progress, and depends on the support of the Secure Stomp Messenger[20] protocol, which is being integrated and tested.

[18] https://wiki.egi.eu/wiki/Globus_integration_task_force

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

[20] https://wiki.egi.eu/wiki/APEL/SSM

Two new integration activities started in PQ7 concerning QosCosGrid[21] software and desktop grids. QosCosGrid middleware is needed for advance reservation of compute resources, a feature needed for tightly coupled usage of EGI and PRACE resources. The QosCosGrid (QCG) middleware is an integrated platform offering advanced job and resource management capabilities to deliver to end-users supercomputer-like performance and structure. Integration activities are being carried out in collaboration with the MAPPER project[22]. QCG integration with GOCDB and SAM are almost completely accomplished. Desktop grid integration is a shared effort of EGI and EDGI and aims at the seamless integration of desktop infrastructures into EGI operations. Desktop grid integration into GOCDB and SAM Update 17 are finalized, while accounting integration is still in progress.

Another major technical obstacle which concerns integration is the possibility to harvest information about deployed resources and their status from the Information Discovery System regardless of the middleware stack supported. A series of workshops were organized involving EGI, EMI and IGE to collect requirements and define a shared implementation roadmap.

2.2.1.1.3 Help Desk & Support Teams

EGI Helpdesk. A big step towards permanently keeping up to date the documentation about the Support Units (SUs) connected to GGUS was taken. All GGUS-FAQs about SUs were updated and migrated from PDF documents stored by the GGUS server to the central EGI wiki[23] and are now fully searchable. Various legacy SUs were decommissioned, and new SUs for VO support, operations support and 3rd level support, were created.

The Technology Helpdesk has been established to provide specialized software support. Through the established interface to the EGI-RT system the first EMI-1 release and subsequent updates were announced to EGI via the Technology Helpdesk.

In order to properly support monitoring of SLAs established with external technology providers, and to monitor status and progress of NGI support services, requirements were gathered from within the project and from external projects (EMI and WLCG) for the implementation of a report generator. This is expected to become operational in PY3. A prototype was presented at the EGICF 2012.

Usability of GGUS was enhanced and new features were introduced: availability of Resource Centre downtime information when submitting a ticket (in order to avoid unnecessary incident notifications if the affected Resource Centre is under maintenance) , extended search capabilities by user DN and type of problem, usage of "Problem Type" values when submitting TEAM or ALARM tickets, issuing of warnings in case of ticket updates, the automated switching after a deadline of tickets in status "solved" or "unsolved" to “verified”, the addition of new fields in the ticket the middleware product affected.

[21] http://www.qoscosgrid.org/trac/qcg

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

[23] https://wiki.egi.eu/wiki/GGUS:FAQ

GGUS was interfaced to Service NOW, CERN helpdesk system. As many EMI 3rd level middleware support units are located at CERN, processes were defined and implemented for ticket synchronization between the two systems.

GGUS was updated roughly on a monthly basis. Several Remedy updates were deployed and auto-switching between different web front-ends was implemented for an improved availability of the system.

VO Support. EGI VO Services aim at supporting VOs in the whole process of start-up, management and operation, pointing out to tools, services, documentation and guidelines to maximize the usage of the resources, easing service deployment, and bridging the VO community with the infrastructure need. The operations community is in charge of operating VO-specific services (both operational and functional depending on the user needs), and of supporting operations and users through the EGI helpdesk. The VO services are mature enough to be supported by NGI operational teams and the expertise on operating those services is also widely available in the operations community. The infrastructure of VO functional services comprises more than 700 service instances. The overall number of international and national VOs registered in the Operations Portal amounts to 226 (+3.20% yearly increase), including 20883 registered users (+14.30% increase). High-Energy Physics, Astronomy Astrophysics and Astro-particle Physics, and Life Sciences are the mostly active disciplines with respectively 93.6%, 2.25% and 1.30% of the overall EGI used normalized CPU time in PY2.

VO SAM, VO Admin Dashboard, and LFCBrowseSE are now mature systems supporting VO operations and being deployed by interested NGIs and/or VOs to assist them in VO daily operations and management. The first prototype of the VO Operations Portal[24] – released by JRA1 and fully integrated into the Operations Portal – was deployed and feedback was provided before finally rolling it to production. Handling of operations issues raised by VOs and VRCs was streamlined, and these are now regularly discussed during OMB meetings. A new procedure for VO decommissioning was drafted and it will be finalized by the OMB during PY3.

Grid oversight. Central Grid Oversight (COD) of EGI was responsible of the certification of new NGIs being created either as a result of legacy EGEE federated operations centres stopping operations, or because of new Resource Providers joining the infrastructure.

COD has been producing newsletters[25] on a monthly basis to improve the dissemination of technical information to NGI operations support teams. A change to the COD escalation procedure was discussed and approved at the OMB for the streamlining of oversight activities (procedure PROC01[26]). A set of tutorial videos was prepared.

As to the follow-up of under-performing sites, COD investigated the issue of the high rate of UNKNOWN monitoring results for some NGIs, which could undermine the meaningfulness of the

[24] https://operations-portal.egi.eu/voDashboard

[25] https://documents.egi.eu/document/298

[26] https://wiki.egi.eu/wiki/PROC01

availability/reliability reports. Frequent UNKNOWN results are usually an indication of an unstable/mis-configured RC or of unreliable local monitoring infrastructures. A technical analysis of the problem was conducted and NGIs have been contacted monthly if they have too high a percentage. The issues were successfully resolved by the majority of NGIs.

COD also participated to the definition of the specifications of a Nagios test that will automate the notification of performance issues to sites, so that administrators are proactively warned and can take counter measures during the course of the month. The Nagios probe measures availability and computes it daily across the last 30 calendar days. It returns WARNING if 70% ≤ availability ≤ 75%, and CRITICAL if availability <70%. A prototype version (provided by SRCE) will be available in March for testing. This work will be completed in PY3.

A new performance indicator (called "ROD performance index") [27] was prototyped and now is rolled to production to measure NGI support performance on a monthly basis. The ROD performance index is an indicator of the number of tickets and alarms that the NGI had trouble to handle in due course (it is the sum of the number of daily tickets "expired" and of daily pending alarms older than 72 hours). NGIs with ROD performance index larger than 10 are supported to improve their on duty activities. This task is contributed to reduce the index, and a gradual decrease was observed during the months it was prototyped. As of October 2011 NGIs not meeting the minimum performance threshold are requested to provide explanations in GGUS tickets.

As of January 2012 COD is also responsible of follow-up of NGI core services that are under performing, and of providing support and documentation.

Network support. Network support activities are responsible for disseminating and supporting tools for network monitoring and trouble shooting, and for IPv6 testing of deployed software.

HINTS is the tool for the execution of on-demand tests and measurements to facilitate troubleshooting network problems. Testing of HINTS was carried out in France and in Italy and probes were installed in various sites. A development server was set up in Paris and a production one is located in Rome.

Network support sessions were organized and co-located with the EGITF 2011 and EGICF 2012.

A task force on IPv6 compliance testing of the deployed technologies (ARC. gLite, GLOBUS, UNICORE) was kicked off. The implementation of the IPv6 testbed is in progress. Activities are run collaboration with the IPv6 HEPiX working group.

2.2.1.1.4 Grid Management

Migration of EGI central tools from gridops.org domain to egi.eu was successfully finalized in PQ5. Decommission of the legacy gridops.org domain was carefully planned in order to no compromise the functionality of the operational tools and of external systems depending on the tool end-points. All EGI central tools are now using egi.eu domain.

Operations Portal.

[27] https://wiki.egi.eu/wiki/Grid_operations_oversight/ROD_performance_index

Seven upgrades were deployed in production during PY2. New major components of the portal were released: the VO Dashboard and the Security Dashboard, and VO management features were greatly enhanced. The Operations Portal infrastructure comprehends to date one central instance and four NGI instances (NGI_BY, NGI_CZ, NGI_GRNET and NGI_IBERGRID). The old Operations Portal (cic.egi.eu) was decommissioned in PQ8.

Service Availability Monitoring

Seven different SAM releases[28] were deployed in the distributed infrastructure following the EGI software release procedure. The SAM infrastructure comprises 26 NGI instances serving 35 EGI partners, 2 federated instances serving 4 the Russian federation and the Asia Pacific federation, one central instance for monitoring of the EGI.eu operations tools and 3 external ROC instances for Canada, IGALC and Latin America.

The installation of a secondary instance is now possible for the implementation of a high availability configuration. This is an important feature as SAM is critical for the reliable collection of monitoring results and the computation of availability statistics. Starting with SAM-Update13 SAM uses a new test to check the EGI Trust Anchor versions on worker nodes. The new test is included in OPERATIONS tests and AVAILABILITY tests (and consequently has an impact in case of failure on the Availability/Reliability monthly reports). The main new feature of the new CA test is that metadata provided in the CA release is used so that there is no need for manual updates of the CA probe package for each new CA release.

GOCDB

A new release was rolled to production (v. 4.2) supporting scoping of Sites and Service Endpoints into EGI and Local categories (Sites and Service Endpoints marked as being part of the infrastructure are exposed to the central operational tools while local entities are not considered part of EGI. This release also rolled into production many bug fixes and a large scale re-factoring of the database as part of the earlier v4.1 release.

The set-up of the failover configuration of the master instance was completed. This includes a 2 hourly export and refresh of the secure download for the database, and testing of the DNS switching mechanism. The secondary instance is hosted by the Fraunhofer-Institut für Techno- und Wirtschaftsmathematik.

Version 4.3 was released on April 18th. Major change in this version was introduction of the support of service groups (previously known as virtual sites) and new roles.

GOCDB documentation on wiki was greatly improved.

Accounting

The migration from R-GMA to messaging across the entire distributed infrastructure was successfully completed. A new ActiveMQ STOMP consumer was deployed in pre-production for external testing; it is relevant to all the accounting systems who publish summary records directly into the APEL database (NGI_IT, NGI_NDGF, OSG and a few Resource Centres).

[28] https://tomtools.cern.ch/confluence/display/SAMDOC/Release+Notes

The APEL team engaged with developers from EMI (for ARC and UNICORE), IGE (GLOBUS) and EDGI (for Desktop Grids, which are under integration) in a series of phone calls and meetings to discuss how they will publish accounting data to the central repository. A TCB task force on accounting[29] started in February to coordinate deployment and requirements gathering across different projects and infrastructures.

A new accounting portal release "Canopus" (v4.0) was deployed in production. The new version uses a new codebase and versioning system, with GGUS referenced commits. URL management was improved and the pchart 2.0 graph engine was introduced. In addition, FQAN-based accounting views were implemented for "VO manager" and "site" views.

Messaging

During PY2 the broker network underwent major software upgrades to improve reliability, scalability and operability. These will continue in PY3. The new version of the messaging broker ActiveMQ 5.5 was tested in October 2011 and subsequently rolled to production. For testing purposes an additional broker network was set up. The testing network consists of 4 brokers (2 at AUTH and one at CERN and SRCE).

Purpose of the upgrades was the rolling into production of new features and the improvement of the messaging infrastructure in various ways. Reliability and availability of the messaging system was enhanced through the usage of virtual destinations. Scalability was improved to reduce the number of connections to the broker network that are left pending, and the implementation of a test network was completed to try new software releases. The difference between "camel routes" and "virtual destinations" is in how data is consumed. With camel routes a message is recorded until it is consumed and then deleted, while with topics a message is published to a consumer without keeping record. A time to live of 3 days is adopted by default. This improves the reliability of message delivery. ActiveMQ 5.5 also supports dynamic failover.

Metrics Portal

The first production release of the Metrics portal[30] was open to the public and used for the recording SA1 metrics from QR6. New metrics, HTML and Excel reports for NGIs have been developed in response to user needs.

Availability

The current OPS availability profile used for the computation of Resource Centre availability and reliability statistics, was discussed and approved for extensions. CREAM, ARC-CE and lcg-CE monitoring results, are now included in the computations. In collaboration with WLCG the replacement of GridView with the Availability Calculation Engine (ACE) was approved.

From May 2011 a new suspension policy for under-preforming sites was introduced that increases the limit of the availability for Resource Centre suspension from 50% to 70%.

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

[30] http://metrics.egi.eu/

The EGI Operations Level Agreement framework was considerably extended in PY2. The Resource Centre Operational Level Agreement[31] (OLA) v. 1.0 was finalized and approved in May, and was subsequently updated in PQ8. The OLA task force[32] was extended to start working on the Resource Provider OLA: the first Resource infrastructure Provider Operational Level Agreement[33] was drafted and approved in PQ6. This is a major accomplishment as for the first time a document defining NGI responsibilities and services was approved. The OLA will be incrementally expanded as monitoring and Availability/Reliability reporting evolve, to include additional NGI services. The OLA structure was also reviewed in order to conform to the ITIL[34] best practices.

Following the approval of the OLA, the first NGI Core service Availability/Reliability report was distributed starting in September 2011, and as of January 2012 NGIs are requested to provide a service improvement plan in case of under-performance. The reporting framework was extended in order to extract results from the MyEGI portal and produce summarized Availability/Reliability reports. These for the moment only include top-BDII.

The EGI sites availability re-calculation procedure was finalized (PROC10[35]) was approved and subsequently updated in PQ8. To support this procedure, and in general to support Resource Centres and NGIs in case of problems with the distributed performance reports, the Service Level Management Support Unit was created in GGUS.

Core services

The DTEAM VO assists RC administrators and operations teams in troubleshooting. Its support is a mandatory requirement for all RCs that deploy VO-enabled middleware. It is served by two geographically distributed VOMS servers in Thessaloniki and Athens. During PY2 seven new DTEAM NGI groups were created and 3 ROC Groups were decommissioned (ROC_Italy, SEE and DECH). The DTEAM VO is successfully used by a large number of NGIs, and AUTH – the partner responsible of DTEAM VO management – was responsive and supportive.

The EGI Catch-All Certification Authority is an important service for new user communities and to support user authentication in the early stages of creation of a new grid infrastructure. It is currently serving 5 countries which do not have a national accredited Certification Authority (Albania, Azerbaijan, Bosnia and Herzegovina, Georgia and Senegal).

A TOP-BDII, a WMS and an LB service were installed as catch all services for NGIs that do not operate their own services for the site certification process (especially for small NGIs). In addition a

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

[32] https://wiki.egi.eu/wiki/Task_forces#Operational_Level_Agreement_.28OLA.29

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

[34] ITIL ® is a Registered Community Trade Mark of the Office of Government Commerce, and is Registered in the U.S. Patent and Trademark Office.

[35] https://wiki.egi.eu/wiki/PROC10

portal was built, that synchronizes with GOCDB and gives the ability to the NGI Managers to add and remove on demand uncertified sites from the catch-all TOP-BDII.

Documentation

The EGEE portal for documentation (GOCWIKI) was phased out in September 2011, after a phase of update and migration to the EGI wiki of relevant pages. The transfer of material to the EGI wiki is now complete. The best practices manual is now complete and a new best practice manual on the management of the top-BDII in failover configuration was approved.

Three new procedures were approved: Re-computation of monitoring results and availability statistics (PROC10[36]), Resource Centre Decommissioning Procedure (PROC11[37]), and Production Service Decommissioning Procedure (PROC12[38]). The remaining ones were periodically updated as needed.

In addition, documentation, training and support wiki pages were significantly updated by EGI.eu:

- https://wiki.egi.eu/wiki/Operations_Manuals
- https://wiki.egi.eu/wiki/TrainingGuides
- https://wiki.egi.eu/wiki/Support

2.2.1.1.5 Tools

2.2.1.1.5.1 Operations Portal

The development of the Operation Portal during the entire PY2 was focused on the following components

1. **The security dashboard** allows for collecting results from Pakiti[39] and Nagios[40] in order to identify security issues, to display on a NGI basis the relevant sites with identified security problems, to create tickets against sites, to transmit results so as to prevent leakage, to compute security metrics based on the numbers gathered

2. **The VO management module and the VO ID card**: the VO management module allows for managing the VO ID card structure and contents and for checking the VO ID cards registration and update using a persistent VO IDcard repository. The VO Management module is now used by COO (VO management is now operated by SA1).

[36] https://wiki.egi.eu/wiki/PROC10

[37] https://wiki.egi.eu/wiki/PROC11

[38] https://wiki.egi.eu/wiki/PROC12

[39] http://pakiti.sourceforge.net/

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

3. **The VO oriented dashboard** allows for showing aggregated status information of sites from test performed using VO Nagios systems. The implementation of this module was one of the main new features approved by the OTAG during PY2 for this tool.

A first prototype of **the security dashboard** was delivered in July 2011 and after evaluation from the EGI CSIRT group[41] it was refined during PQ6, in particular for what concerns authentication (based on GOCDB and EGI SSO to allow for restricted access to the data according to scope-based policies), visualization and metrics generation (customisable, dynamically generated and available as chart or tables in various formats). Various iterations with the CSIRT for the evaluation allowed improving the tool on each Operation portal release.

Updates to the **VO module** addressed a number of small requirements coming from operations and user support teams. In particular to provide more reliable statistics on users numbers and VOMS server information (to validate a VOMS host a "Test URL" button has been added to directly query the "listMembers").

The **VO oriented dashboard** development was approved at the OTAG-10 held during EGITF2011 and a first prototype was made available for evaluation in February 2012 and improved to address community comments in March and April.

The maintenance and requirements addressing for other components (i.e. Dashboard, COD Dashboard, Broadcast tool) continued using the available effort.

The **Regional Packaging** of the portal was maintained and the generation of the package (and the related documentation) is now done automatically. The synchronization with the central instance has been improved to allow special characters in the messages. The broadcast module has been added in the regional version. There are currently four instances of the regionalized operations portal used by: Portugal/Spain, Greece, Belarus and Czech Republic. The NGI_UK will set-up a regional instance in 2012. The package and the documentation are available through a SVN server[42].

Decommission of the former CIC Portal has taken place now that all of its features have been successfully migrated to the new Operations Portal. The last remaining feature (the User tracking tool) was migrated during PQ5, but some users were still relying on the old xml feed for the VO information available in the old instance, consequently the decommissioning will finally take place in PQ8.

During PY2 a total of 10 releases were performed by the Operation portal team, detailed release notes are available[43].

2.2.1.1.5.2 EGI Helpdesk

There have been several activities in PY2:

1)The continuous **integration of NGIs** into the support infrastructure. Thirty-three NGIs support units were created and integrated into GGUS by the end of PY2.

[41] https://wiki.egi.eu/wiki/Csirt

[42] https://cvs.in2p3.fr/operations-portal/regional/trunk/

[43] http://operations-portal.egi.eu/aboutportal/releaseNotesBrowser

2)The refinement of the **Technology Helpdesk.** This is a module of the main tool used by the project to handle all middleware related issues, and acts as a bridge between EGI and its external technology providers. It was a major development task during PY2. Currently two workflows are handled in the Technology Helpdesk: (i) the software support workflow for bugs discovered in production; and (ii), the software provisioning workflow. In the latter case, this starts with a release announcement by one of the technology providers, and ends with the acceptance or rejection of the proposed software products.

3)The development work to have an **active-active fail-over system** for the data, the logic and the presentation layers of GGUS will continue in PY3.

During the first half of PY2 many requirements were created for the GGUS **report generator**, a module that allows to create statistics about tickets, support units and reaction times, so during the OTAG-10 meeting it was agreed to organise a workshop to analyse the requirements and propose a unique formalised set of requested features that could be implemented taking into account the available effort for GGUS within JRA1. The workshop was held on October, 26th and 27th at KIT in Karlsruhe, Germany and the agreed new features were:

- SLA reporting for EMI, IGE & SAGA
- Calculation of response times and solution times
- Calculation of statistical values like median, average, maximum, minimum of the data collected
- Reports by different parameters like priority, status, support unit, site etc.
- Ticket lifetime for user's view

A first prototype of the report generator was presented at the EGICF2012 and the first usable version will be made available in June 2012 to have it finalised by the end of the year.

The GGUS product team published a total of nine releases during PY2.

2.2.1.1.5.3 GOCDB

During PY2 effort focused on the three main activity streams: data model and backend, User Interface, Deployment. In the following some details are given about the developments for each of the mainstreams.

1)Data Model and Backend

a.Data scoping in the central GOCDB implemented in GOCDB-4.2, was released 25-11-2011. The business logic, user interface and PI queries were updated to include new 'scope' tags used to label entities as belonging to the 'EGI' or 'Local' groups (similar in concept to a tag cloud). For more information see section 4.2 regarding a Regional-Publishing GOCDB and https://wiki.egi.eu/wiki/GOCDB/Release4/Regionalisation.

b.Service Groups (also known as Virtual Sites), which form logical groupings of services that are potentially distributed across a number of different sites and regions. See https://wiki.egi.eu/w/images/f/fe/VirtualSitesDesign.pdf.

c.Refactorization of the backend PROM database to enable transaction demarcation in higher level business routines. Released in v4.1 on 01-11-2011 (v4.0 could leave the db in an inconsistent state if an error occurred since full rollback over lower level functions was not atomic)

d.Update of wiki and documentation that now provides detailed information about ongoing developments, release notes and user guides

e.New services and middleware types integration such as QoSCoSGrid[44] middleware and iRODs[45] services

f.Certification Status changes recorded in new audit/history table with accompanying PI queries

g.Improved coding abstractions and integration of established design patterns and unit tests

2)User Interface

a.An improved user interface and a new Model View Controller logic

b.Integration of a new XML output module based on the Query2XML package to generate complex nested XML documents from GOCDB data. This was required to render the Service Groups in the new 'get_service_group' PI method

c.A new finer grained roles and permissions model. Released April 10th 2012. More information available at https://wiki.egi.eu/wiki/GOCDB/Release4/Development/NewRoles

d.Increased frontend responsiveness using db connection pooling

3)Deployment

a.A detailed versioning schema that was added to the tool releases starting from GOCDB v4.1 on 01-11-2011

b.Established a GOCDB failover instance (http://bit.ly/frau_gocdb). The failover imports a secure download of the GOCDB data every 2 hours. A manually triggered DNS switch for the

[44] http://www.qoscosgrid.org

[45] https://www.irods.org

'goc.egi.eu' domain between the production server and the failover server is in place. During normal operation, the failover is read only in order to prevent data-synchronization problems.

2.2.1.1.5.4 Accounting Repository

Focus of the accounting development team during PY2 was, besides the maintenance of the production quality tool, on the migration of the system to a new transport method based on the Secure Stomp Messenger (SSM) which uses the EGI Messaging Infrastructure. SSM transfers files in a secure and reliable manner between sites. SSM acts as a producer and consumer, taking files from a source directory at the publishing site, signing them, encrypting them, and sending them as messages to the receiving site where they are unencrypted and placed in a target directory. The necessary handshaking is also done using the messaging infrastructure to confirm the delivery.

A test service for the new STOMP/Python publisher/consumer was implemented in PQ5. Discussions with a variety of stakeholders including those that publish by direct database insert were started and continued during the entire PQ6. A training workshop was organised during the EGITF2011 where the current accounting architecture and the planned roadmap were presented[46].

In PQ7 and PQ8 the APEL team completed the APEL Accounting Repository major redesign and the development for the new transport layer (SSM) was completed.

Accounting for new resource types

The activity for TJRA1.4 during PQ5 & PQ6 focused on the investigation of what is currently available for the accounting of new resource types, identifying overlaps and possible collaborations in the work performed by various stakeholders, trying to bring many of them together. The types of accounting considered, and the stakeholders were:

- CPU (inc OGF UR) (EGI, EMI, OGF)
- MPI (EGI, EMI)
- Storage (inc StAR) (EMI, OGF, EGI)
- Virtualisation (EGI, other projects)
- Applications (EGI)
- Data Use (EUDAT, PaNData)

The agenda[47] and minutes[48] of the workshop at the EGITF2011 provide further details.

The following paragraphs describe the activity performed on clouds, parallel jobs and storage accounting.

The JRA1 accounting product team is also participating to the following task forces:

[46] https://indico.egi.eu/indico/sessionDisplay.py?sessionId=71&confId=452#20110922

[47] https://www.egi.eu/indico/sessionDisplay.py?sessionId=89&confId=452#20110922

[48] https://www.egi.eu/indico/getFile.py/access?contribId=361&sessionId=89&resId=0&materialId=minutes&confId=452

- Integration of new middleware types (i.e. UNICORE[49])
- Integrated Accounting[50]

to understand the requirements about the integration of new middleware stacks (ARC, UNICORE, gLite, Globus).

Accounting for clouds

The JRA1 accounting product team actively participated in the EGI Federated Clouds taskforce[51] and led the so-called "Scenario4- accounting" activities with the mandate to include into the accounting system records coming from the federated cloud infrastructure. A "straw man" Usage Record based on the current CPU record has been defined and the partners running the various infrastructures are reporting back on how easily this can be integrated with the local cloud infrastructure. Prototype testing with the cloud Usage Record following started before the EGICF2012 and formed part of the demo shown at the event. It will be improved in PY3 and will be more reliable once the new production infrastructure based on SSM is in place

Accounting for parallel jobs

The JRA1 accounting product team also participated in the definition of the accounting for parallel jobs usage record within standardization bodies (i.e. OGF UR-WG). It was confirmed that it was sufficient for parallel jobs to publish at least the number of compute nodes used by a job[52], the need for publishing also the number of physical CPUs is under debate. It is intended that CPU elements in a UR record contain the number of cores for compatibility with the accepted practise for serial jobs running on a single core. Unfortunately, APEL does not support these fields and this will be a target for PY3. The accounting repository will be ready to receive record of this type once the database is migrated but will not actually receive the data from most sites until the new APEL client is rolled out.

Storage Accounting

The JRA1 product team also participated in the definition of the accounting storage record within standardization bodies (i.e. OGF UR-WG). Meanwhile it is moving forward as an interim standard in EMI where the storage product teams are preparing to populate records. They will publish using SSM and the APEL repository will receive the records and load them into a database. The release of the EMI products is not scheduled until EMI-3 in May 2013 but this may be negotiable if EGI has a strong requirement

2.2.1.1.5.5 Accounting Portal

During PY2 the focus of the Accounting Portal team was on the redesign of the tool that now adopts a new codebase, a new graph engine and a new interface towards the Accounting Repository that now is

[49] https://wiki.egi.eu/wiki/UNICORE_integration_task_force

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

[51] https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce

[52] https://forge.gridforum.org/sf/wiki/do/viewPage/projects.ur-wg/wiki/ConsideredURFieldsCPU

based on EGI's Message Infrastructure. A query optimization was performed in order to increase the performance. Many bug fixes and some new features were released to production. Among the new released features there are:

- Implementation of an XML interface to obtain data from the Custom View as requested by the user community

- Change of the Operations Portal interface from Oracle queries to XML (needed because he old Oracle interface was abandoned - the VO manager and user information would have stopped working)

- FQAN data on "VO admin" and "Site Admin" views as requested by the user community

The release process was also improved on a six months release base as described in the Accounting Portal Roadmap.[53]

2.2.1.1.5.6 SAM

A total of seven SAM updates were released to production during PY2. All updates went through the full EGI Software Provisioning process[54] as was decided at the beginning of the project. Another release was prepared and is ready to be released. The main activity streamlines on which the development focused are: 1) the inclusion of new middleware types probes 2) the improvement of the user interface for myEGI 3) the creation of a profile management system (POEM) to ease the addition of custom probes to the framework.

Concerning the inclusion of new middleware probes, now the SAM system contains probes for gLite, ARC, UNICORE and GLOBUS. Desktopgrid probes are under integration. Other important probes that were included/improved in the system are those for GLEXEC, for the Certification Authorities and for other EGI operational tools such as GOCDB, GGUS, Accounting Portal and Metrics Portal. Moreover some of the internal SAM components can now be monitored (i.e. the MRS).

A first integrated version of the Profile Management system (POEM) was released as a prototype. The new profile management system introduced a major rework of the existing SAM architecture as it completely changed the bootstrapping of all the major components and introduced significant changes to the overall data flow.

The MyEGI interface was improved during the year and many features were added, in particular all the features of the old gridmap[55] tool are now available in myEGI including the possibility to show the summary of Logical CPUs, Physical CPUs and HSPEC06 numbers for a given site.

Another important feature that was added in PQ6 is the possibility to configure a failover instance in a hot stand-by mode[56]. To ease the packaging and distribution of the system a clean-up of dependencies and metapackages reorganisation were also performed. Decommissioning of the entire old SAM

[53] https://documents.egi.eu/document/517

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

[55] https://openlab-mu-internal.web.cern.ch/.../EDS_GridMap_Poster.pdf

infrastructure including Gridview, Gridmap and the old SAM database and Web interface was finalized in April 2012 following successful deployment and transition to the new SAM/Nagios infrastructure.

Message Broker Configuration Support

The focus for this activity during PY2 was on improving the authentication and authorization system to access the EGI production broker. Various updates of the ActiveMQ version were also performed. A system to enforce idle connection eviction was put in place in order to greatly improve the performance of the system. A plan for the implementation of a credential-sharing procedure between brokers that are members of the same network was also created. This is needed in order to have the same credentials available to all brokers and allow clients to choose on their own which endpoint to use. This task has to be finished before requesting clients to migrate to authenticated connections.

A lot of effort was spent to support migration from "topic" to "queues" (that enables the message history retention) of other operational tools, in particular the Operations Portal, allowing for a more reliable message exchange between the tools.

2.2.1.1.5.7 Metrics Portal

The metrics portal was completely redesigned with respect to the one available in the EGEE-III project, but progress was slow in PY1 due to hiring problems at CESGA. The definition of the roadmap and development plan started at the end of PY1 after a series of meetings with the project activity leaders, held to get the list of requirements that has since driven developments.[57]

At the beginning of PY2 effort was focused on the development of connectors to GGUS, BDII, and availability metrics data sources. The development portal was cloned to the production one. The possibility to manually override automatic metrics was added. In PQ6 the authorization model was redesigned according to the provided requirements and various enhancements were added (i.e. aggregated metrics on NGIs). The following features were added during PQ7:

- Per country metrics (only on requested NGIs).
- Heavy query optimization.
- XLS output support.
- Many metric fixes and modifications.

A second major update (v2) for the portal was released in PQ8 enhancing all the portal components

[56] https://tomtools.cern.ch/confluence/display/SAMDOC/Installing+SAM+Update-13#InstallingSAMUpdate-13-FailoverNagiosconfigurablehotstandbymode

[57] https://documents.egi.eu/document/516

2.2.2 User Support

2.2.2.1 Shared Services and Tools

2.2.2.1.1 Dashboards

Monitoring of the distributed infrastructure and the activities of the user communities on this infrastructure is a vital for ensuring its quality and performance. Monitoring is of particular importance for Heavy User Communities (including HEP) due to the scale of their activities and the quantity of resources that they are using. There are two main tasks that have to be addressed by the monitoring systems used by HEP VOs: monitoring of the distributed sites and services, and the monitoring of the VO activities, namely job processing and data transfer. The Experiment Dashboard was developed in order to address the monitoring needs of the LHC community, but in contrast to other monitoring systems it provides common solutions that work transparently across various middleware platforms and are not coupled with VO-specific frameworks, offering instead a common way to instrument those frameworks for publishing monitoring data.

During PY2 the Experiment Dashboard applications were heavily used by the LHC virtual organizations (VOs), in particular by ATLAS and CMS. The system plays an important role for everyday operations, for site commissioning activity and for the distributed computing shifts. More than 200 CMS physicists daily access CMS Dashboard task monitoring in order to follow processing of their tasks on the distributed infrastructure. ATLAS DDM Dashboard is being actively used for monitoring of ATLAS Data transfers. All LHC experiments use Dashboard SAM portal which provides information for evaluating site usability from the VO perspective.

The performance, scalability and functionality of the system are steadily improving following the growing scale of the LHC computing activities and the requirements of the LHC community.

Most of the Dashboard cluster had migrated to the virtual machines. It facilitated deployment and maintenance tasks.

All Dashboard applications use ORACLE for persistency. In the beginning of 2012, ORACLE cluster at CERN which hosts Dashboard DBs migrated to new hardware and ORACLE 11g. This resulted in substantial improvement of the performance of the Dashboard applications.

One of the important goals of the Experiment Dashboard development is to offer common monitoring solutions for the LHC VOs. Achieving this goal permits the monitoring infrastructure to be simplified and contributes to its sustainability. All Dashboard development in 2011 and beginning of 2012 followed this strategy. During the referenced period, the global transfer monitoring system which provides cross-VO monitoring of all data transfers on the WLCG infrastructure was prototyped and deployed on the preproduction server. New versions of the job monitoring applications were developed based on the common hBrowser framework. hBrowser framework offers common client-side implementation of the user interfaces with generic look and feel and with the advanced plotting, sorting and filtering functionality. After successful validation by the user community, new versions of the job monitoring applications were deployed in production. The Site Status Board (SSB) which provides a common solution for monitoring of the distributed sites and services was redesigned in order to improve performance of the data queries and the quality of the user interface. New SSB

versions were successfully validated by the LHC VOs, were deployed in production and received a positive feedback of the user community.

The LHC VOs rely on the Service Availability Monitor (SAM) for remote testing of the distributed sites and services. Therefore one of the important WLCG monitoring tasks during 2011 and beginning of 2012 was migration to a new SAM architecture. The Experiment Dashboard provides a portal for navigating of the SAM test results and site availability calculated. The new portal compatible with the new SAM architecture was validated and deployed in production.

Further development of the Experiment Dashboard system is being aligned with the LHC computing needs and will follow the requests of the LHC virtual organizations.

2.2.2.1.2 Tools

Ganga is an end-user tool for creating and managing computational tasks. The mature Python codebase provides for a stable yet extensible framework, which is consistently used by over 400 end-users per month, distributed over 100 individual sites and submitting in excess of 250,000 jobs per week. It remains a popular environment for running grid analysis jobs within the LHCb and ATLAS experiments, and is known to be used by a further 10 communities across a range of scientific disciplines.

During the past year development centered around migration of community-specific applications into the Ganga core, from where they can be exploited by a wider user base. For example, the GangaATLAS Tasks package, which was initially developed to allow bulk submission of Athena jobs, is now available as an experiment-neutral tool. This new implementation extends the existing functionality to allow job-chaining workflows and automated (though rate-limited) resubmission of failed jobs. Further development has exposed the Task functionality to LHCb users, who are now able to create book-keeping query objects to return a specified set of input data files.

Similarly, the 'prepare' method, which was also originally only available in the GangaATLAS package, has been extended to a range of other applications. This provides the ability to lock an application in a known configured state, such that it can be rerun at a date in the future. In contrast to the original ATLAS implementation, the new approach stores all files required to run the application in the user's Ganga repository (c.f. the system temporary directory) thereby ensuring their persistence. Thus, the manual book-keeping and file-copying that was previously required to share configured applications between jobs is no longer required.

In parallel to the above, evolution of the existing Ganga packages has continued, highlights include:

●Improvements in the GangaPanda package to improve the handling of output files. Specifically, the ability to merge multiple output files has been added and job splitting mechanisms have been re-implemented, alongside resubmission routines for failed or killed subjobs.

●In the area of tag-file support both user-generated tag files and centrally produced tag datasets are now fully supported within Ganga. The tag-based workflow has been integrated into the ATLAS PanDA distributed analysis system and provides the foundations on which the ATLAS Skim service runs.

●The ProdTrans application was developed for submitting jobs to ATLAS production queues from within the HammerCloud test framework; providing support for running production transformation jobs.

●Significant improvements were made to the Ganga testing framework, its associated tests and web frontend.

●The release procedure tools were improved to provide a greater level of automation and to allow some elements to run in parallel. This resulted in a reduction in the time taken to generate new Ganga software releases.

●Improvements were made to the end-user documentation; the internal Python document strings were re-written or expanded for some modules and a feature was added that prints the latest release notes when a new version of Ganga is executed for the first time.

Current development within the core is focused on providing generic output-handling routines which allow output to be compressed before being shipped back to the user or, alternatively, sent to remote storage technologies. The resulting dataset can then be viewed and downloaded directly from the Ganga interface.

Future developments will centre around those aimed at increasing the adoption of Ganga as an analysis platform. For example, inclusion of Ganga within a recognised Linux distribution would provide both a familiar installation procedure for new users, and reinforce the fact that Ganga is a mature and established tool.

2.2.2.1.3 Services

2.2.2.1.3.1 Grid Relational Catalog (GRelC)

The GRelC service is a Grid database management service aiming at providing access and management functionalities related to relational and non-relational databases in a Grid environment.

During the EGI-InSPIRE project, the GRelC software will support the HUC communities with a new set of functionalities available as web application through the GRelC Portal and the DashboardDB interface[58]. The DashboardDB will represent the proper web access interface to the EGI Database of Databases. Such a "registry" will complement the functionalities provided by the EGI Application Database and will represent a distributed and multi-VO system supporting the HUC. It will be searchable and will allow people for instance to post comments and share opinions about the EGI data sources, ask to get access to a specific Grid-DB and join discussion groups.

A major achievement in PY2 was the first production release of the *DashboardDB registry*. It has been deployed and integrated into the official GRelC website. The functionalities tested, validated and now available to the end users in the official release are: user registration, user profile management, Grid-database registry, Grid-database search and discovery, discussion groups (including group subscription and notification), Grid-database rating and tagging. Moreover the *permalink* feature allows the registry to be easily "exported" (as a gadget) in other web contexts (the AppDB could be an example in the near future). This functionality has been exploited to integrate the DashboardDB registry into the

[58] https://documents.egi.eu/document/746

official GRelC website, under the tab “Registry”. The back-end of the DashboardDB (developed during the first year) is also running in the production release.

Another important feature released in this period, is the *DashboardDB global monitoring* view which allows the monitoring of the GRelC instances deployed in EGI. Such a monitoring part complements the existing monitoring tools already available in the EGI environment, but provides a different perspective (which was missing before) focused more on “Grid-database management” services. By design and to address re-usability, also the DashboardDB global monitoring view is available as a permalink. This means it could be easily "exported like a gadget" in different web-contexts representing an add-on to more general EGI monitoring tools.

Significant progress was made in producing a *gLite 3.2 compliant release of the GRelC middleware*. The porting is now complete and the new RPMS are available on the GRelC website along with the necessary documentation (both for end-users and administrators). An integration plan of the gLite 3.2 compliant release of the GRelC software into the Italian Grid Initiative (IGI) release has been defined and started. A new repository at the IGI level has been setup for this task.

Strong support has been provided to the HUC (in particular to ES and LS).

A use case (UNIPROT) related to the LS community has been jointly defined with the bioinformatics group working at the University of Salento and implemented over the past period. The UNIPROT data bank (about 13GBs flat file, now moved into 30 relational tables) is now available in Grid to this community through the GRelC service interface. A set of tools have been also designed and developed to support this use case and the management of the UNIPROT data bank. The porting of the UNIPROT data bank is an advantage because in a structured database the users are able to carry out complex queries that they were not able to perform on the flat file. Basically they are now able to run join queries across multiple information/concepts. Redundancy is also removed by definition, considering the intrinsic normalization process in building the relational data bank. Having the data bank into a relational database allows users to fully exploit the power of a declarative language like SQL. A key point is that we are porting several flat-file data banks into relational databases and this will allow researchers to cross-reference multiple (federated) data banks at the same time by running SQL queries.

The end users profiting from this are bioinformatics. In particular a group at the University of Salento is testing it for its own research activities. There are other potential users that could be interested in exploiting the new data bank and we'll contact them in the coming weeks, once we'll get some preliminary input from our beta-testers. It is up to the end users to decide which version (relational or flat-file based) of the UNIPROT data bank they want to exploit for their purposes. If they are interested in extracting information by running complex join queries they would probably use the relational DB (as it is happening at the University of Salento), whereas if they want to run some existing data mining applications they will probably exploit the flat file. Of course, the two data banks (relational and flat-file based) can co-exist in the grid infrastructure.

Porting/adapting or developing new tools is up to the end users. SPACI is providing the ETL software to port the flat-file data banks into relational databases. It is comprised of a C library responsible for extracting the data from the flat file and for storing it into an XML file and the GRelC translator component responsible for translating the XML file entries into bulk queries for the new relational

database. Last but not least, the access to the UNIPROT data bank can take advantage of the features provided by the GRelC grid database interface (several kinds of delivery mechanisms, fine grained security control, etc.).

Another use case (IAS, *Invasive Alien Species*) in the Life Sciences domain and concerning biodiversity aspects has been defined in the last period jointly with the Biology group at the University of Salento taking into account their needs and requirements. This use case foresees the adoption of the GRelC middleware to support a platform for sharing invasive species information at a global level in the *Italian LifeWatch Virtual Laboratory*. Even though the use case is starting considering a single GRelC service instance at the University of Salento in Lecce (Italy) with some biodiversity databases, the long term plan (1 year) foresees more than ten nodes across the Italian country and hosting several databases exposed in Grid through the GRelC Grid-database interface. This interoperability exercise of data stored in distributed databases is meant to map the fragility of ecosystems to alien species invasion and estimating a long-term impact of alien species on biodiversity.

Concerning the Environmental domain, the GRelC service is exploited as a Grid metadata service at the Euro-Mediterranean Centre for Climate Change – (CMCC, Italy) to manage experiments/datasets metadata (compliant with ISO geospatial standards). Specific support for XML documents has been provided due to the complexity of the metadata hierarchy managed at the CMCC.

Visibility has been addressed updating the GRelC website and the wiki, giving oral presentations and tutorials, writing new papers and organizing sessions and workshops in ICT (e.g. International Conference on Computational Science - ICCS2011) and domain-related conferences (e.g. European Geosciences Union - EGU2011) as documented in the QRs.

The support in terms of management, monitoring and control of the GRelC services provided through the DashboardDB will be further extended and improved. The DashboardDB will provide new (service-specific) views about the GRelC deployment and the status of the registered services. A strong dissemination about the DashboardDB registry will be carried out in PY3 to register the GRelC service instances already deployed in EGI as well as the Grid-database resources they offer. Moreover, the questionnaire prepared during PY1 will be improved, extended and disseminated among the HUC to attract new users and make the existing ones aware of the new Grid-database registry tool. A roadmap towards the *European Middleware Initiative Release* will also be defined.

Since the DashboardDB registry has been designed as a self-consistent component (it can be easily exported and embedded in other web contexts like a web gadgets) it will be publicized as a new *EGI gadget* at the European level (www.egi.eu/user-support/gadgets).

Additional support to the HUC will be provided to address new user needs and requirements. In particular the UNIPROT and IAS in the Life Sciences domain will be further supported. A major goal regarding the support for the LS community will be to create in Lecce (SPACI node) "*a Grid database node hosting several data banks addressing LS needs*". Concentrating at the same place many different biological databases will be relevant for the LS community and could be also crucial to attract new users.

Finally website, training events, tutorials, talks and papers will represent additional ways to disseminate the results of this activity.

The GRelC roadmap towards the European Middleware Initiative release can be considered an important milestone for the next period. Anyway, it can be considered a long term task since it will need continuous support to be implemented over the next years.

The DashboardDB registry as a whole can represent another long-term task due to the relevant role the registry can play in a production Grid environment. Moreover, even though it now contains Grid-database entries related to the network of GRelC services, in the future it could also address the needs coming from other tools in the same research area, becoming a more general and comprehensive Grid-database registry service.

For more information please refer to the GReIC section of the HUC Software Roadmap:

2.2.2.1.3.2 SOMA2

SOMA2[59] is a versatile modeling environment for computational drug discovery and molecular modeling. SOMA2 is operated through a WWW-browser and it offers an easy access to third-party scientific applications. The SOMA2 environment offers a full scale modeling environment from inputting molecular data to visualization and analysis of the results, and including a possibility to combine different applications into automatically processed application workflows.

During PQ5 CSC maintained and operated CSC's SOMA2 service. In addition, Autodock 4 integration work in SOMA2 continued. Also, we started investigations of how to setup a SOMA2 service which would also be provided to other user communities and not just for current CSC users. Program development effort was put in to upgrading the core UI library components and migration work was needed to be able to use the updated components. Work was also done to improve the basic Grid support in SOMA2 along with other minor fixes and improvements.

During PQ6 the focus was on a working Autodock 4 integration in SOMA2. Currently this provides a virtual screening service facilitating distributed resources. In addition, a new release of SOMA2, including the Grid support and more, was prepared. Basic setup for the SOMA2 service, which would be provided to other user communities as well and not only for current CSC users, also proceeded. However, this was slowed down by SOMA2 release not being ready yet. Also, CSC maintained and operated CSC's SOMA2 service.

The main outcome of PQ7 was a new version release of SOMA2. On 31.1.2012 SOMA2 version 1.4 (Aluminum) was released and made available on the SOMA2 web site. This release contains most of the EGI related development work so far consisting of the SOMA2 Grid integration and other improvements. Most of the work performed in PQ7 was put to finalizing and packaging the version release. Also during PQ7, CSC maintained and operated CSC's SOMA2 service.

The outlook for 2012 based on the achievements so far; we will enable a Grid-enabled application service for interested user communities. At the moment it seems that CSC is going to set up a separate pilot service for this. We will also operate and support the existing SOMA2 services. In addition, we will advertise the Grid enabled SOMA2 service to different user communities.

The long-term goal is to expand the scientific applications integrated in to the SOMA2 service, and integrate application services from different Grids hosted by different virtual organizations into

[59] http://www.csc.fi/soma

SOMA2. This should be set as an important milestone because from the end users point of view, this would make using scientific applications in different Grids very easy and transparent. The SOMA2 service will continue to be maintained and operated, while seeking new scientific applications to be added as part of the service. The development of SOMA2 gateway will continue according to feature roadmap including possible feature requests from the user community, bug fixes and other enhancements.

2.2.2.1.4 Workflow & Schedulers

During PY2 the work with Kepler (and the Serpens suite for Kepler) has been focused on developing new template use cases basing on user's requirements, improving the performance of the Kepler actors, and performing the tests of the different use cases. In addition thanks to the dissemination activities performed during the project, the template use cases initially prepared for the Fusion community have been customised and reused by users from other communities (including astrophysics and computational chemistry).

The first astrophysics workflow has been also customised and tested. It controls the production of realistic simulations of the anisotropies of the Cosmic Microwave Background (CMB) radiation. These simulations will be used to test the performance of various algorithms designed to detect the presence or absence of non-Gaussian features in the simulated CMB maps, before they are applied to actual data as observed by Planck, a satellite from the European Space Agency. In order to test the algorithms we need to produce large numbers of simulations. Each one of them is made of a combination of a Gaussian and non-Gaussian component plus realistic instrumental noise that takes into account the observing strategy of the satellite. The initial workflow follows the parameter study workflow template.

Also the first workflows from the CompChem area have been arranged and customised. This workflow follows the template of having application chain with the parameter study use case and the collector included. This use case is divided into two phases. First, the GFIT3C application fits the user provided ab-initio points of the potential energy surface. This is executed as a standalone task. The result is passed to the ABC, the quantum reactive scattering application, with a set of multiple alternative input files to be examined.

The dissemination activities carried out in different locations and meetings have allowed the impact of this work to be shown to other communities. The improvements introduced in the actors to improve reliability and scalability, can be of interest for many other communities and can potentially extend the usage of this tool.

2.2.2.1.5 MPI

The MPI sub-task produces numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry, Earth Sciences, Fusion and Astronomy and Astrophysics (A&A) communities, and beyond. In addition, it focuses on ensuring that the user communities and site administrators benefit from several rudimentary improvements to the methodologies used and the available documentation. Many of these objectives are iterative, often requiring updates or fine-tuning. Other objectives, such as participation at the EGICF and the EGITF, will be repeated at regular intervals. The core sub-task objectives (which bring definition to the tasks sustainability) are:

- Improved end-user documentation, addressing MPI application development and job submission in ARC, gLite and UNICORE;
- Quality controlled MPI site deployment documentation, to date over 120 sites in EGI support MPI;
- Outreach and dissemination at major EGI events and workshops;
- User community, NGI and site engagement, gathering direct input;
- Participation in selected standardisation bodies.

During PY2, TCD and UNIPG started independent work on exploiting GPGPU enabled devices. The proposed models have different implementations; however both focus on using hybrid grid/cloud solutions. UNIPG worked on grid implementations of many parallel codes, including CHIMERE, RWAVEPR, ABC, DL_POLY, NAMD and VENU96. Many of these codes and their associated algorithms have direct re-use by other scientific domains.

CSIC has contributed to the SIESTA computation chemistry application software. CSIC has made continuous improvements to the MPI documentation wiki and has also tested the MPI-START integration with SGE and PBS like job managers with newer versions of MPICH2. CSIC is investigating incorporating support for MVAPICH into the MPI-START framework. MVAPICH has advanced Infiniband and fault-tolerance support.

TCD produced a proof-of-concept work using MPI-START to support generic parallel workloads. This will be demonstrated at the EGICF Parallel Computing Workshop. TCD will also demonstrate simple OpenMPI/OpenCL hybrid code running on the Grid.

The MPI SA3 team participates in the MPI-Virtual team. The formation of the team has help increase the visibility of the MPI task within the user communities, and a dedicated GGUS support unit has also been established to support sites and users. This has had a significantly good impact in ensuring that MPI related tickets are assigned to the proper expert team.

2.2.2.2 Domain Specific Support

2.2.2.2.1 High Energy Physics

The Services for High Energy Physics (HEP) task continues to focus primarily but not exclusively on the 4 LHC VOs (experiments) centred at CERN: ALICE, ATLAS, CMS and LHCb (and hence for and via the WLCG project and collaboration). Services and tools developed or extended by these VOs are also used by other HEP experiments and/or are under consideration for the future.

The work of SA3 goes hand in hand with that of the WLCG project in general: what is described below specifically describes the contribution of SA3 however it should be considered in the broader context of the entire project by considering its role within the WLCG project as well.

Not only were SA3-supported tools and services an important cornerstone of these achievements but the drive for common solutions made further inroads during this period – an impressive result given that stability is the norm during data taking.

During the past year WLCG launched a number of Technical Evolution Groups (TEGs) in which EGI-InSPIRE SA3 members (as well as others from EGI-InSPIRE) played key roles. The final report from

these groups is not expected until May 2012 and is expected to drive the further evolution of WLCG and hence have an impact on grid sites and projects involved.

A brief overview of these TEGs, the involvement of EGI-InSPIRE, as well as the main recommendations of each group are given below.

2.2.2.2.1.1 Analysis Tools and Support

2.2.2.2.1.1.1 HammerCloud

HammerCloud (HC)[60] is a Grid site testing service developed around Ganga. HC uses frequent short jobs to validate a site's availability and functionality, and also delivers on-demand stress tests to aid in site commissioning or general benchmarking. HC was developed with the ATLAS experiment but is used today by CMS and LHCb experiments in addition.

The main achievement of the past year was a merging of different codebases between the three VO users to one generic codebase, which is customizable per VO and will allow extensions to new VOs in the future. CMS and LHCb were already running a common codebase at the start of the year, whereas ATLAS was running on older independent code. During the ATLAS instance migration to the generic codebase, a number of performance issues were discovered; some sub-optimal database queries were found and the VO-agnostic model layer was found to hurt overall performance of the web interface and backend logic.

After some Django optimizations were applied (including adding related-table hints), performance was restored (in some cases the number of database queries required to generate a page view decreased by 3 orders of magnitude.)

A second large achievement in the past year was an extension of the "job robot" functionalities to validate the VO Grid sites. This HC testing has been seen as very successful by the ATLAS VO; improvements to the auto-exclusion service led to ~25% decrease in Grid failures for ATLAS users. This motivated additional testing of the Monte Carlo Production Grid queues, this work was developed during the past year and put into production in mid-February. With this achievement, HC is now validating all ATLAS Grid sites and job workflows. For CMS, the job robot functionalities have been progressively improved over the past ~6 months the tests are regularly queried to understand the site reliabilities. Feedback from these tests presently depends on human operators, in future it is hoped that some automatic feedback into the experiment's workload management systems will be possible. LHCb has been working in the last quarter to integrate the HC testing with the LHCbDIRAC resource status system, this work is still in progress.

One of the plans for the past year was to find a more stable database service to host the HC databases (previously the HC team has been running its own MySQL servers). It was originally planned to port the HC object layer to Oracle, however a new hosted MySQL service became available in CERN IT during PY2, so the databases have been migrated to this service. The plan to use Oracle has therefore been scratched. Moving forward, the plans for 2012 include:

- Integration of LHCb testing with the LHCbDIRAC resource status system;

[60] https://savannah.cern.ch/projects/hammercloud/

- Working on a standard benchmark (or benchmarking suite) to evaluate the storage elements (e.g. standard set of ROOT test jobs and datasets);
- Investigation of error and performance correlations across VOs (at multi-VO sites);
- Development of a multi-VO testing interface for users to easily test all experiment workflows;
- HC installation simplification and packaging for other-VO usage.

Regarding organization and sustainability, the HC project is currently led by an EGI-InSPIRE funded FTE, and developments and operations are complemented by at least 1 partial FTE from each user VO (ATLAS has 2 extra support persons (partial)). It is foreseen that the continued life of the project does not depend one any one of these partners.

2.2.2.2.1.1.2 CMS CRAB Client

The CMS Remote Analysis Builder (CRAB) was the first analysis tool in CMS to aid users in configuring CMS applications for distributed use, by discovering the location of remote datasets and submitting jobs to the Grid infrastructure. CRAB has progressed from a limited initial prototype nearly 5 years ago to a fully validated system that is critical to the production of published physics results. CMS currently observes more than 400 unique users submitting CRAB jobs per week, with close to 1000 individuals per month. The CMS Computing Technical Design Report (CTDR) estimated roughly 100k Grid submissions per day. In the second half of 2010 the job submissions routinely exceeded the estimate by 40-50% and CRAB coped well with the increased load.

After the development of the CRAB submission client a CRAB server was developed, which has increased the scalability of submission and added capabilities of automatic resubmission. The CRAB server also provides a development platform for additional capabilities.

During the past year effort has been spent both to maintain the production version of CRAB and to develop & commissioning the new one. Particularly the maintenance of CRAB version 2 has focused on:

- Fixing problems with JobSplitter as well as deprecating srmv1[61] protocol;
- Guaranteeing the compatibility of the newer version of the CMS Analysis framework (CMSSW) and the underlying condor and gLite middleware changes;
- Supporting the CERN analysis facility migration to the EOS system.

Regarding the CRAB3 development and commissioning activities the main areas of work were:

- Development of the RESTFul based CRABServer interface. APIs for basic functionality have been implemented; a long test and integration phase has been managed. Thanks to the feedback received over the past year further optimization of the REST APIs have been addressed both to improve the functional and technical aspects. In particular, all the security related implementations needed for the final production version are now ready. The validation of the whole set of inputs provided to CRAB by the user is included in the upcoming version;
- Implementation of the User Sandbox system. The User Sandbox is the archive of all user own library, files, configurations which must be shipped to the worker node. A central cache HTTPS based

[61] https://www.gridpp.ac.uk/wiki/SRM

system is now available;

- The new crab client has been made available, supporting all the functionality associated to the resources exposed by the REST layer. The client is now implementing the support for the CRAB2 to CRAB3 configuration file translation. Once moved to CRAB3, the user configuration file will be only python based with the previous text based approach to be removed;
- The first version of the software with the AsyncStageOut functionality, which asynchronously moves the user produced output, has been tested, integrated and released. A basic specific monitoring has been implemented to support the related operations;
- Both the user and code documentation for all the CRAB3 components have been developed using sphinx (http://sphinx.pocoo.org/);
- The first version of the deploy/manage system has been implemented. It is based on the RPM, which indeed has been produced for all the single components of the new CRAB system;
- A lot of effort was spent to support the first multi-user integration test. This was done in the middle of October and the scope of the test is to collect user feedback needed to consolidate the tool and decide the priorities in order to move from development to production quality;
- Several cycles of four days face-to-face meeting have been organized in order to include discussion of both follow-up to the beta tests and development and scheduling of specific items.

The main items of development for the next year are:

- To rewrite the whole monitoring implementation currently developed, moving from a drill-down to a push model approach. The aim is to build a centralized monitoring service, based on CouchDB, which is able to collect monitoring documents from the distributed agents;
- To implement advanced functionality like user data publication, support for the user output merging, and all what will come-up from the users feedback;
- To evolve the current client to a python library. The aim is to enable the possibility to be used by external bot/services like HammerCloud, which will be the first case we will support;
- To evaluate, eventually integrate, and use frontier on worker node to cache UserFileCache;
- To perform a large scale test, including the whole stack which is in the CRAB3 system, namely: Client, REST Interface, UserFileCache, WorkQueue, AsyncStageOut and distributed Agents.

2.2.2.2.1.2 Data Management Tools and Support

2.2.2.2.1.2.1 ATLAS Distributed Data Management Tool

ATLAS, one of the LHC experiments, fully relies on the use of Grid computing for offline processing and analysis. This processing is done worldwide using the well-known tier model across heterogeneous interoperable Grids and the ATLAS Distributed Data Management (DDM) project is responsible for the replication, access and bookkeeping of ATLAS data across more than 100 distributed Grid sites.

The work during the last year has focused on the consolidation of the current production system in different fronts:

- The DDM infrastructure monitoring is responsible for the early detection of system failures and degradations of DDM services. The infrastructure has been extended to a client-server model that communicates over message queues and where the server publishes the health reports to the CERN IT Service Level Status monitoring solution. Based on this new implementation, significant effort has been invested in developing and improving the service monitoring for a variety of the DDM subcomponents.
- Optimization of the DDM Site Services, which is the set of agents responsible for the ATLAS data placement using the underlying EGI middleware (mainly FTS, LFC and SRM). Improvements for this component include:

  ○Optional submission of FTS jobs to GridFTP (bypassing SRM). This was required for the improvement of the data import and export on the new EOS storage element at CERN.

  ○Adaptation of Site Services for Tier3s in order to provide a throttled way of transferring data to these sites. The requirements of these sites are:

    ▪Submission of FTS jobs to the GridFTP endpoints.

    ▪No file/dataset registration in the LFC and Central Catalogues: Sites want to have local control of the data without taking care of catalog synchronization.

    ▪Since there is no LFC registration, file look-up has to be done using *uberftp-ls*.

  ○Adaptation to use the recently implemented ATLAS Grid Information System.

  ○There have been several releases with general enhancing, fine-tuning and bug fixing.

    ▪Optimized usage of temporary storage areas for cross-cloud transfers.

    ▪Pinning of source replicas to ensure their availability.

    ▪Implemented new policies for tape staging according to latest operational experiences.

  ○Migration of the Site Services instance serving US sites from BNL to CERN.

  ○Evaluation of running the service on virtual machines with lower CPU and memory profiles.

- Support has been provided for a variety of DDM components, with strong focus on the day-to-day operations in data replication.

The current ATLAS DDM software is now in a mature state and the present work is focused on maintenance and support operations. Inside the ATLAS Distributed Computing community there is an on-going discussion about the proposal of the CERN-PH-ADP group of developing a new DDM system (the Rucio project) to solve the current shortcomings and scalability issues in the Central Catalogues. The details about the future of the project and the implications in other groups are unclear at this point.

2.2.2.2.1.2.2 CMS Data Management Tool

Building on the previous experience acquired by the ATLAS experiment, the CMS Popularity Service has been developed to monitor the experiment's data access patterns (i.e. frequency of data access, ac-

cess protocols, data tiers, users, sites, CPU usage). The understanding of this framework provides a crucial step ahead towards the automation of data cleaning and data placement.

In addition, a fully automated popularity-based site-cleaning agent has been deployed in order to scan the Tier2 sites that are reaching their space quotas and suggest obsolete, unused data that can be safely deleted without disrupting analysis activity. The implementation of this agent is based on the initial ATLAS code, which has been re-factorized to a plug-in architecture with a common core for the CMS and ATLAS experiments. The system has been significantly improved by adding a database backend that stores the suggested cleaning and a new web frontend.

Future work in this area has mainly two points of focus:

- Firstly, CMS is currently exploring how to extend the Popularity system with the usage patterns of the data accessed through the XRootD protocol.
- Secondly, CMS wants to analyse the possibilities of dynamic data placement strategies and demonstrate that the current Workload Management infrastructure could call out to perform data replication of hot data to facilitate its analysis. This would require interactions:
    - With the Data Management and Workload Management systems for the hook to replicate the data and decide to which sites the replications should be done.
    - With the Popularity Service and dashboard for automatically making additional replica of hot data and to not make additional replicas at sites that currently are oversubscribed.
    - Interactions with the Data Distribution system (PhEDEx) to discover available resources and make the replication requests.

2.2.2.2.1.2.3 LHCb DIRAC

DIRAC framework provides a complete solution for using the distributed computing resources of the LHCb experiment. DIRAC is a framework for data processing and analysis, including workload management, data management, monitoring and accounting[62]. LHCbDIRAC framework is the DIRAC extension specific to the LHCb experiment, which has been formally separated from DIRAC in order to streamline the implementation of features requested by LHCb community.

Continuing the work of PY2, in PY3 the work will cover:

- Consolidation of the service for data consistency checks between file catalogues and SEs, especially in the operational aspects. In particular, procedures should be established for notifying the VO's data management team about the files which were found in an inconsistent state, and decide about their removal. This will streamline the task of consistency checks and will enable the data manager to reduce the time needed to fulfill the task.
- Support and possible improvement of the accounting plots of storage resources usage developed and put in production during the past year. This system, already extensively used by members of the collaboration, could undergo some improvement, taking advantage of new features of the LHCb DIRAC DMS that are planned to be implemented during next months.
- Adapt when necessary the DIRAC DMS to changes in the data management middleware (in

[62]https://documents.egi.eu/document/540

particular SRM and LFC).

- Continue the development of the popularity service, just started during PQ7. Such a service should provide metrics to asset the data-sets popularity and provide a ranking of the most popular data-sets (i.e. data most frequently accessed by users). The final goal is to use the information about data popularity to implement a dynamic data placement model, where the number of replicas of a given data-set is related to its popularity. This would considerably help the VO's data managers to optimize the usage of storage resources on the Grid and would automate many operations that are currently done manually and thus time expansive.
- Provide general support for LHCb computing operations on the Grid, both for production and private user's activity.

2.2.2.2.2 Life Science

The Life Science Grid Community (LSGC) was set up during PY1 with the central role to the coordination and support of the Life Science activity on EGI infrastructure. In the context of the EGI-InSPIRE project, the Life Science HUC contributed to the LSGC effort by maintaining a production quality Grid environment for Life Sciences by providing technical skills and manpower for VRC operation, as well as some specific tools dedicated to the Life Science community. A significant effort is spent on the VO-wise monitoring and trouble-shooting of the EGI infrastructure. This effort is split between a daily trouble-solving activity to ensure the immediate usability of the Life Science resources for the community, and a longer term effort in VRC management tools that are being developed to simplify and lighten the VRC administrators' workload in the future. Moreover, the provision of additional services, mostly a Grid database interface (GRelC) and a data encryption service (Hydra) are being handled within the HUC.

2.2.2.2.2.1 LSGC User Management Tools

The Life Sciences HUC steers the LSGC ("Life Sciences Grid Community" VRC[63]) effort to organize the community and deliver new services. A particular effort is invested in assisting users to better exploit the Grid and rationalizing Grid usage. In this context, several services to the HUC users have been provisioned:

- Web gadgets listing Life Sciences applications and community requirements posted to the RT systems set up by User Community Support Team have been added to the LSGC wiki;
- A HUC support service is delivered. A technical team of expert users has been set up. It addresses the difficulties reported by users on the VRC mailing lists or through GGUS. Bi-monthly phone conferences are scheduled and shifts are organized to ensure that there is always a team on duty tackling the problems. See technical team wiki for details[64];
- The technical team pro-actively monitors the infrastructure health at a VO level, to identify the problems occurring. The French NGI deploys a dedicated Nagios server[65] for that purpose. New probes to monitor all VO SEs, WMSs and CEs were developed;

[63] LSGC wiki, http://wiki.healthgrid.org/LSVRC:Index

[64] Biomed technical team wiki, http://wiki.healthgrid.org/Biomed-Shifts:Index

- On-line reporting tools easing the monitoring of SEs space management have been added to the technical team wiki page.

Thanks to the tools set up, a large-scale action to clean up all obsolete files let behind by years of Grid usage from a heterogeneous community with 200+ registered users is currently being conducted.

Recently accomplished work:

PQ5:

- Setting-up of community communication channels and technical assistance to the end users.
- Provisioning of a redundant VOMS server for the biomed VO.
- Investigation of technical means to similarly deliver a redundant LFC server.
- Strengthening of standard procedures to deal with storage resources decommissioning or failures.

PQ6:

- Dissemination of good practices as well as knowledge of tools used / developed among the VRC technical team.
- Improvements of the Nagios monitoring probes.

PQ7:

- Monitoring of new resources, in particular all WMSs and CEs accessible to the biomed VO users.
- Publication of a paper[66] describing the role, the organization and the impact of the Life Sciences technical team in experiments support.

This work is fully integrated with EGI. The LS community does not intend to deploy an independent parallel infrastructure.

The development of a HUC users database and management tools to assist VO administrators in their daily task is planned. This service will interface to Virtual Organization Membership Service (VOMS) servers as well as the EGI applications database, to avoid replicating existing information. It will complement the VOMS and applications database with extra-information on the users and their affiliations. It will be interfaced to the LSGC Dashboard.

2.2.2.2.2.2 GRelC data access interface

The Grid Relational Catalog (GRelC) access interface is a general purpose Grid database service that can be exploited in several Life Sciences use cases.

In this regard, two new use cases have been investigated during PY1. The first one concerns with the UNIPROT gene data bank (which has been completely imported into a new relational database and made available through the GRelC interface) whereas the second one (which is still being conducted) is a biodiversity use case, aiming at studying and sharing invasive species information. A detailed description about the two use cases is available in Section 7.

[65] Biomed Nagios server, https://grid04.lal.in2p3.fr/nagios

[66] "Technical support for Life Sciences communities on a production grid infrastructure", HealthGrid 2012. http://hal.archives-ouvertes.fr/hal-00677839

Recently accomplished work:

PQ5:

- Availability of the UNIPROT proteins data bank through the GRelC interface.

PQ6:

- Improvement of the ETL (Extraction-Transformation-Loading) tool. The ETL tool is comprised of a C library responsible for extracting the data from the flat file and for storing it into an XML file and the GRelC translator component responsible for translating the XML file entries into bulk queries for the new relational database.

PQ7:

- Definition of a new use case on biodiversity.

The support to the Life Sciences domain will be a key activity also in PY3. New use cases will be defined jointly with LS groups and the existing ones (e.g. the one on biodiversity at the University of Salento) will be refined and finalized. A more detailed and cross-domain outlook about the GRelC service is discussed in Section 7.

2.2.2.2.2.3 Hydra encryption service

Hydra[67] is a file encryption/decryption tool developed by EMI to enable the protection of sensitive files stored on Grid storage resources. The service is composed by a distributed encryption key store (hence its name), and client command lines that can (i) upload/fetch keys to/from the key store and (ii) encrypt/decrypt data files using these keys.

An experimental Hydra service has been successfully deployed since the end of PQ4 (April 2011), first on a 32bits gLite version 3.1 and then on a 64bits gLite version 3.2 server. There are two difficulties identified for delivering a production service using the Hydra key store. First, many production sites are mis-configured, having deployed older version of Hydra in a former gLite release and sometime exposing hydra tags that are not valid any more or older version of the hydra client. Second, the current version of the hydra client is developed as part of the EMI middleware and its dependencies are incompatible with the gLite 3.2 release deployed on the production infrastructure. The service delivered today is therefore only a test service mostly used for the validation of the functionality delivered.

Recently accomplished work:

PQ5:

- Deployment of a prototype Hydra service on a gLite release 3.1 UI;
- Work on the migration to gLite release 3.2.

PQ6:

- Provision of first beta-release of the Hydra server software compatible with gLite release 3.2;
- Fixes in the installation and configuration procedures.

[67] Hydra service overview. https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS

PQ7:

- Work on Hydra client provision for sites providing resources to Life Sciences;
- Site configuration problems are being tackled. Work on Hydra CLI package with EMI to try installing it on the current release of gLite deployed on the production infrastructure.

To deliver a production service, Hydra depends upon:

- Server's provision for Hydra key stores;
- Standard client packages installation on all EGI computing resources;
- Service monitoring and maintenance.

For the future; a 3-server based Hydra key-store will be deployed as a service for the life sciences community. This task, due in the first year of the EGI-InSPIRE project, has been delayed until gLite 3.2 (64 bits) servers could be provisioned and further slowed down by the end of the HealthGrid association which was supposed to host one of the servers. An attempt will be made to produce client packages compatible with the current middleware release and deploy them. If it proves impossible, it will be needed to wait for the first deployment of the EMI release on the EGI infrastructure. The LSGC technical team will handle the service client installation and monitoring for more sustainability. The service monitoring procedure will be integrated in the Nagios box-based monitoring service offered by the LSGC technical team.

2.2.2.2.3 Astronomy and Astrophysics

Activities carried out by the A&A community during PY2 have been focused on the following topics:

- Coordination of the A&A community focusing in particular on the long-term sustainability plan;
- Visualization tools and Services: VisIVO;
- HPC, parallel programming, and GPU computing;
- Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure.

2.2.2.2.3.1 VisIVO

During PY2 significant results have been achieved for what concerns the porting to the Grid of VisIVO[68] (Visualization Interface for the Virtual Observatory), a visualization and analysis software for astrophysical data. It consists in a suite of software tools aimed at creating customized views of 3D renderings from many types of datasets.

The activity mainly focused on the study and on the porting of the VisIVO MPI version to a gLite Grid. The relevance of this activity can be easily understood if one considers that, depending on the structure and size of datasets, the Importer and Filters components could take several hours of CPU to create customized views, and the production of a movie could take several days. For this reason the MPI parallelized version of VisIVO plays a fundamental role.

The preliminary study and the porting were mainly focused on the most important modifications of the code necessary to make VisIVO MPI fully compliant with the gLite Grid. The VisIVO MPI version, in fact, works assuming that the shared home directory and each process can directly work on the tables.

[68] http://visivo.oact.inaf.it/index.php

Some classes of the VisIVO Filter component were modified to allow selections on a data table and preliminary tests were carried out.

Another important aspect that would enhance the performances is the integration of VisIVO on Grid nodes where GPUs (Graphics Processing Units) are available. GPUs are emerging as important computing resources in Astronomy as they can be successfully used to effectively carry out data reduction and analysis. The option of using GPU computing resources offered by Grid sites to make visualization processing on VisIVO was then considered.

CUDA (Compute Unified Device Architecture) is the computing engine available in NVIDIA GPUs and accessible to software developers through variants of industry standard programming languages. Because VisIVO is developed in C++, the environment of CUDA is used to develop some computing-intensive modules of VisIVO. This activity started at the beginning of PY2 as a preliminary study on how to produce a CUDA-enabled version of VisIVO for gLite. This first study focused on the porting and optimization of the data transfer between the CPU and GPUs on worker nodes where a GPU is available. To provide a service able to take advantage of GPUs on the Grid, A&A acquired a new system (funded by the Astrophysical Observatory of Catania). The server is configured as a Grid computing node.

The heaviest VisIVO Filter, the Multi-Layer Resolution Filter, was also analyzed. This filter makes possible the inspection of a very large user file (hundreds of gigabytes) to create data for the visualization of an entire dataset with different levels of resolution: starting from a fixed position, that represents the centre of the inner sphere, concentric spheres are considered. Different levels of randomization can be given, creating a more detailed table in the inner sphere and less detailed tables in the outer regions, or vice versa. The region external to the last sphere represents the background.

The performances of some additional VisIVO visualization filters such as randomizer, cut, select and swap operations on huge user data tables were also considered.

Now the first grid-enabled version of VisIVO service has been deployed. It is based on a specific grid-enabled library that allows users to interact with Grid computing and storage resources. The current version of VisIVO is also able to interface with and use the gLite Grid Catalogue.

Although VisIVO has been conceived and implemented as a visualization tool for astronomy, it has evolved as a generic multi-disciplinary service that can be used by any other community that needs 2D and 3D data visualization.

2.2.2.2.3.2 Grid and HPC

A working environment where Grid resources and HPC resources can be combined and used together is of utmost importance for many astronomical applications. Some typical A&A applications that could improve their performances in such an environment have been identified: FLY[69] (a cosmological code developed at INAF-OACt) and Gadget[70] + Flash[71], the most common cosmological codes in Astrophysics.

[69] http://www.ct.astro.it/fly/

[70] http://www.mpa-garching.mpg.de/galform/gadget/index.shtml

The FLY code (a tree N-Body code) was executed on the gLite Grid. Some problems arise when running with a high number of N-Body particles. One of the most critical bottlenecks, in fact, is represented by the data transfer from the catalogue to the effective worker nodes. Unfortunately, the huge dimension of datasets makes tricky to have multiple replicas of them. The same problem happens during the data production phase, when a list of data files (each of them being tens of gigabytes in size) is produced.

To overcome this problem a workflow is being defined which can be adopted when the code is used in challenging scenarios where data files several tens of gigabytes in size have to be handled.

To test cosmological simulations we are now in the phase of identifying and defining some use-cases; the goal of this activity consists in preparing the environment, the input file and the watchdog procedure to verify the run and the results.

Recently the activities related to Grid and HPC are carried out in close coordination with EGI and with IGI[72] (the Italian NGI) in Italy. A tight coordination with EGI and with NGIs, in fact, is crucial given that collaborations and agreements with PRACE and with other entities that operate and maintain HPC resources (typically supercomputing centers) are very important for this activity. One of the activities recently undertaken within the task TSA3.5 aims to introduce small-size HPC resources in Grid. This implies from one side to install and configure HPC clusters (based on low latency/high throughput networks, HPC libraries and tools, modules and compilers) and from the other side to make the grid middleware aware of these resources. The plan is to verify such small-size HPC resources and the related middleware aware version vs. the most popular cosmological applications mentioned previously, namely FLY and Gadget + Flash.

For this reason, in the next period the activity for what concerns HPC will continue in Italy in close coordination with all those communities involved in IGI that need HPC and MPI to efficiently run their applications.

2.2.2.2.3.3 Access to Databases and interoperability with the VObs

Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure remains one of the hot topics in astronomy. One of the most important related activities deals with the identification of use-cases and test-beds (both applications and complex workflows) that require simultaneous access to astronomical data and to computing resources. The data infrastructure of reference in astronomy is the VObs; end users should be able to access astronomical data in Grid through the VObs standards[73] and launch computational tasks on DCIs. In this context two key issues concern a SSO (Single Sign On) mechanism able to grant access to computing and data resources by means of a single authentication (users do not need to authenticate themselves multiple times) and tools/services to access astronomical databases federated in the VObs from DCIs.

[71] http://www.flash.uchicago.edu/site/

[72] http://www.italiangrid.it/

[73] http://www.ivoa.net/Documents/

For what concerns the way of accessing databases from DCIs, GRelC[74] is one of the tools under evaluation in order to verify its ability to meet the most important requirements of the A&A community.

Another important activity related to the integration of data and computing resources is the work in progress to integrate in Grid the BaSTI[75] (A Bag of Stellar Tracks and Isochrones) Astronomical Database and its feeding FARANEC code. A web portal was developed to make-easier the submission of the FARANEC code to the Grid. The portal allows the user to define a set of parameters and to simulate stellar evolutions in Grid, without worrying about the technical details concerning the underlying Grid Infrastructure. The portal is based on P-GRADE web portal. Since the portal is designed for the submission of an arbitrary job in Grid, it was necessary to do some low level "tricks & hacks" to make it able to fully satisfy our needs. The work aimed at implementing a new version of the portal based on the latest version of P-Grade is now in progress. It will be based on Web Services and this should enable a low level integration with the Virtual Observatory.

2.2.2.2.4 Earth Sciences

Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting, climate change and many others.

The presence of Earth Science in SA3 is centred in the implementation and maintenance of interfaces or tools to provide access to Earth Science specific resources from the Grid, in particular to large data infrastructures; for example resources within the infrastructure of the Ground European Network for Earth Science Inter-operations - Digital Repositories (GENESI-DR), or climate data within the Earth System Grid (ESG). The community is supported independently by organisations and NGIs, and additional effort is put into fostering the community and to provide value-added services around EGI. The Services for Earth Science task covers the implementation of data access scenarios, to permit the utilization of Earth Science data resources in Grid jobs. The work can be divided into two efforts, separate but related in substance: discovery and access through GENESI-DR and ESG.

2.2.2.2.4.1 GENESI-DR

The GENESI-DR infrastructure provides a standardized data discovery interface based on OpenSearch and metadata standards for a federation of data repositories. While in the European project behind it (GENESI-DEC) focuses on a central portal as an interactive entrance point, the usage on EGI requires versatile clients such as a non-interactive, bulk oriented, tool. The work on such a command line client[76] that facilitates the usage in Grid jobs is ongoing. In recent months a data transfer component was included that downloads the bulk of resources found through the discovery process. This is required as the discovered data sets are not accessible through a unified channel, but depending on the data set instead through different access protocols such as HTTP, GridFTP or others. Because of changes to the underlying GENESI infrastructure, the tools had to be adapted. The command-line tool

[74] http://grelc.unile.it/home.php

[75] http://albione.oa-teramo.inaf.it/

[76] https://appdb.egi.eu/?#!p=L2FwcHMvZGV0YWlscz9pZD03MDg

was extended with an ncurses based text user interface for interactive command-line usage over a shell, e.g. in the typical situation of a terminal SSH connection to a gLite User Interface. It now also has more sophisticated data handling and a few bugs were fixed to solve some rare crashes. In addition, a system for configuration over a dot file was added, which enables the user to permanently save and modify some variables like the preferred top request site. Documentation with a description of the functionality and a few standard use cases and examples to help users getting started with the utility has been prepared.

Additionally, a flexible web GUI was designed. It uses the OpenSearch interface provided by a running GI-cat instance (v.i.). The web interface accepts search parameters such as: a keyword, lower and upper time limits, a geographical bounding box (see figure 1) and more. The interface leverages the OpenLayers capabilities for defining the geographical bounding box. The list of results is subsequently being used to generate search queries for the respective datasets such as AIRS, MODIS, GOMOS or GOME (mainly depends on the catalogues that are managed by GI-cat). Mostly, these datasets provide OpenSearch interfaces as well. The user defined search parameters from the first steps are then delegated to these second layer interfaces. The found data is presented to the user, who is able to select the required files. Based on this step, a job description file to be submitted to Grid resources is generated.

2.2.2.2.4.2 Earth Systems Grid

The Earth System Grid Federation (ESGF/ESG) is a distributed infrastructure developed to support CMIP5 (The Coupled Model Intercomparison Project, Phase 5), an internationally co-ordinated set of climate model experiments involving climate model centres from all over the world. Data access within ESGF is provided with two main services: OpeNDAP and GridFTP. A site that hosts these services is called a “Data node”. The team that works on ESG interoperability is developing and testing a scenario based on an application from IPSL which uses CMIP5 data (climate model data stored on the ESG). The MPI code of the application is now running successfully on EGI. An important activity was the implementation of a multi-threaded climate data transfer program to download the data from the ESG data nodes. This smart data transfer tool, named ‘*synchro-data*’[77], facilitates the command line, bulk oriented access to ESG CMIP5 data. The tool can download files from the CMIP5 archive in an easy way, through a list of variables, experiments and ensemble members. The user defines one or many templates that describe the desired data, each of them listing variables, frequencies, experiments and ensemble members. The user separately defines a list of climate models to take into account. Using these templates, the tool explores the ESGF Grid and downloads all the corresponding available files. The program may be run regularly to download the possible new files. This tool has been improved considerably during this period. Major added features include fine grained priorities for transfers, support for CMIP5 "ensembles", an installer, an error watchdog, multi-threaded downloading, incremental downloads through keeping a history, caching to limit the stress on ESGF metadata server, additional actions (delete, cancel, retry), and a statistics module to report about remaining download volume, disk space requirements, etc. One of the team members regularly participates in ESGF developer meetings.

[77] http://dods.ipsl.jussieu.fr/jripsl/synchro-data

A critical point for the interoperability with ESG is the challenge of different authentication schemes for EGI and ESG, as they don't belong to the same federation. This situation, which is also relevant for the usage of the GENESI-DR infrastructure from EGI, was presented at the Workshop on "Federated identity system for scientific collaborations" in CERN on June 8-9 2011. Different solutions have been considered to make the use of ESG data with EGI easier. The software NDG Security Stack (http://ndg-security.ceda.ac.uk/wiki) has been set-up as a testing instance and has been reviewed. It was first developed for the NERC DataGrid and further on extended to be used as the Earth System Grid Federation security module. It includes support for OpenID, X509 and SAML; this security module is frequently used in climate data infrastructure. This instance is to be used for prototyping a solution for the authentication problem, which might be a new Credential Translation Service. The work is done in cooperation with Philip Kershaw (STFC, UK), who is working on the Earth System Grid Federation security model and software. The problem has been communicated to other EGI representatives in the form of a short internal report, describing the situation and possible approaches for a solution. Telephone conferences between representatives of EGI, SA3.6 and Earth System Grid Federation have been carried out, to clarify and discuss the situation and future plans. A prototype modification of MyProxy is currently developed that will issue ESG certificates based on EGI certificate authentication.

While the GENESI-DR infrastructure can be used to discover many data sets of various data centres, and also includes now data from OGC web services, such as WMS instances, the data of the Earth System Grid is not discoverable through it. The search and discovery of ESG data sets is thus a separate issue, which is not a comfortable position for the users of EGI. It would be ideal to have single points of access for the search and discovery of data sets from the different sources, as well as a unified client to transfer the data, for easy inclusion in EGI job scripts or applications. The Earth System Grid, though, offers public access to THREDDS catalogues for the hosted data sets on a given Data Node (e.g. http://cmip-dn.badc.rl.ac.uk/thredds/esgcet/catalog.html for the BADC CMIP5 data). Therefore, solutions to unite the GENESI-DR search with a search capability for these dedicated catalogue instances have been searched. The applied solution is the GI-cat service developed by ESSI-Lab (http://essi-lab.eu/gi-cat), which uses a mediation approach to execute remote searches federating the results and remote catalogue crawling building up a local copy of the catalogues to be searched. As it evaluates and transforms the meta data of its sources to a common internal scheme, it can at the same time offer different catalogue interfaces, allowing existing clients the access through different protocols (such as OGC CS/W, OpenSearch, GeoRSS and OAI-PMH). The software has been deployed at SCAI and the clients are tested against it. Unfortunately, the implementation of the protocol differs slightly, so the clients need to be adapted.

For the future; the EGI Earth Science community is in contact with the FP7 project Virtual Earthquake and seismology Research Community e-science environment in Europe (VERCE), which aims at integrating a service-oriented architecture with an efficient communication layer between the Data and the Grid infrastructures, and HPC. A second novel aspect is the coupling between HTC data analysis and HPC data modelling applications through workflow and data sharing mechanisms. Discussions about effective and goal-oriented exploitation of Grid resources in the projects infrastructure have taken place. The projects members were suggested to register for the ESR general purpose Virtual

Organisation to immediately be able to use computational resources. The setup of a dedicated VERCE VO or VRC is being discussed at present.

2.2.3 Software Provisioning

The major goal for SA2 in PY2 was to put the software provisioning infrastructure into action that was designed and tested during PY1. Starting with the EMI-1 release the SA2 activity relatively quickly found out which parts of the process had room for improvement. As with every process that is introduced, it initially required much more effort to make use of the provided tools and artefacts.

2.2.3.1 Quality Criteria

The Quality Criteria management team has continued following the established process of maintaining Quality Criteria for software deployed within EGI. Published every half-year, each final Quality Criteria release is preceded by peer-reviewed draft sets that allow Technology Providers, Platform Operators and Integrators, and User Communities to ensure that the Quality Criteria accurately capture the needs of consumers. Consequently, two final releases of Quality Criteria were published in PY2[78].

Detailed change logs and release notes for each Quality Criteria revision are publicly available[79]. Each change in the documents is tracked with the source of the change and the criteria it affects. In turn, the affected criteria also include in their description a revision log. Changes to individual Quality Criteria[80] are triggered by feedback given through the Software Provisioning process reports produced for each Product (e.g. see Verification report for L&B[81]), and by monitoring the Software Vulnerability mailing list[82].

With the release of revision 3 of the documents, all UMD Capabilities are covered by Quality Criteria. Further, during the preparation of release 3, the Quality Criteria management team started to reorganise how the Quality Criteria are maintained, and presented to its main users, the Software Verification team organised in TSA2.3. While the individual Quality Criteria are maintained and grouped in documents focusing on a specific type of software quality, the team started to develop a mapping table indicating which specific quality criteria are applicable to which specific software product that is due to be provisioned for release to the production infrastructure. An automated script takes the pool of Quality Criteria and the mapping table as input, and produces one single document that contains all Quality Criteria applicable to a specific software product. These documents are the basis on which Software Verifiers is conducting their work.

[78] EGI Quality Criteria, Revision 2, http://go.egi.eu/qualitycriteria-2, EGI Quality Criteria, Revision 3, http://go.egi.eu/qualitycriteria-3

[79] EGI Quality Criteria overview, https://wiki.egi.eu/wiki/EGI_Quality_Criteria_Dissemination

[80] EGI Quality Criteria definition, https://wiki.egi.eu/wiki/EGI_Quality_Criteria_Definition

[81] Verification Report for EMI L&B, https://documents.egi.eu/document/968

[82] This source of information is not publically available; this is by design so that vulnerabilities can be discussed freely between appointed members while reducing the risk to trigger zero-day exploits.

During PQ8 two incidents in the EGI production infrastructure involved the Quality Criteria team to assess how these incidents could be prevented in the future. Discussed at the 11th TCB meeting[83] several suggestions for the Globus 5.2 incompatibility incident, and the VOMS administration incident are currently developed, and will be incorporated in the next scheduled set of Quality Criteria.

2.2.3.2 Criteria Verification

During PY2 the Criteria Verification team has put into operation all services and processes that were designed and tested during PY1. The complete verification process has been documented in the EGI wiki[84], along with guidelines to new and existing verifiers, and how to use the verification test-bed hosted by CESGA[85]. A cheat sheet[86] was developed and maintained to indicate expertise of individual software verifiers for any particular verified software product. It proved well as a management tool to further reduce the effort spent until the actual functional verification of a software product can commence.

Verifiers keep a record of the actual verification effort for each product update that is taken through the process. This information is aggregated together with other metrics taken from the EGI Software Provisioning infrastructure, and published daily in an Excel sheet[87]. Further provisioning metrics related to Technology Provider service performance are collected and published[88].

During PY2, a total of 104 individual software product versions were verified, and released for deployment on the production infrastructure. Spread over seven UMD releases, an average of 15 product updates per released UMD update were made available for production infrastructure deployment (see Figure 1). The EGI Software Repository[89] provides details of the specific software products that are included in each UMD release.

[83] 11th TCB meeting (F2F), http://go.egi.eu/TCB-11

[84] Software Verification process, https://wiki.egi.eu/wiki/EGI_Quality_Criteria_Verification

[85] Verification test bed HOWTO, https://wiki.egi.eu/wiki/EGI_Verification_Testbed

[86] Verification engineer skill matrix: https://wiki.egi.eu/wiki/EGI_Quality_Criteria_Verification_-_Verification_engineer_skill_matrix

[87] Software Provisioning metrics reports, https://rt.egi.eu/rt/SA2/sa2-sw-rel-verification-metrics.xls

[88] https://rt.egi.eu/rt/SA2/SLA/index.html

[89] EGI Software Repository, http://repository.egi.eu

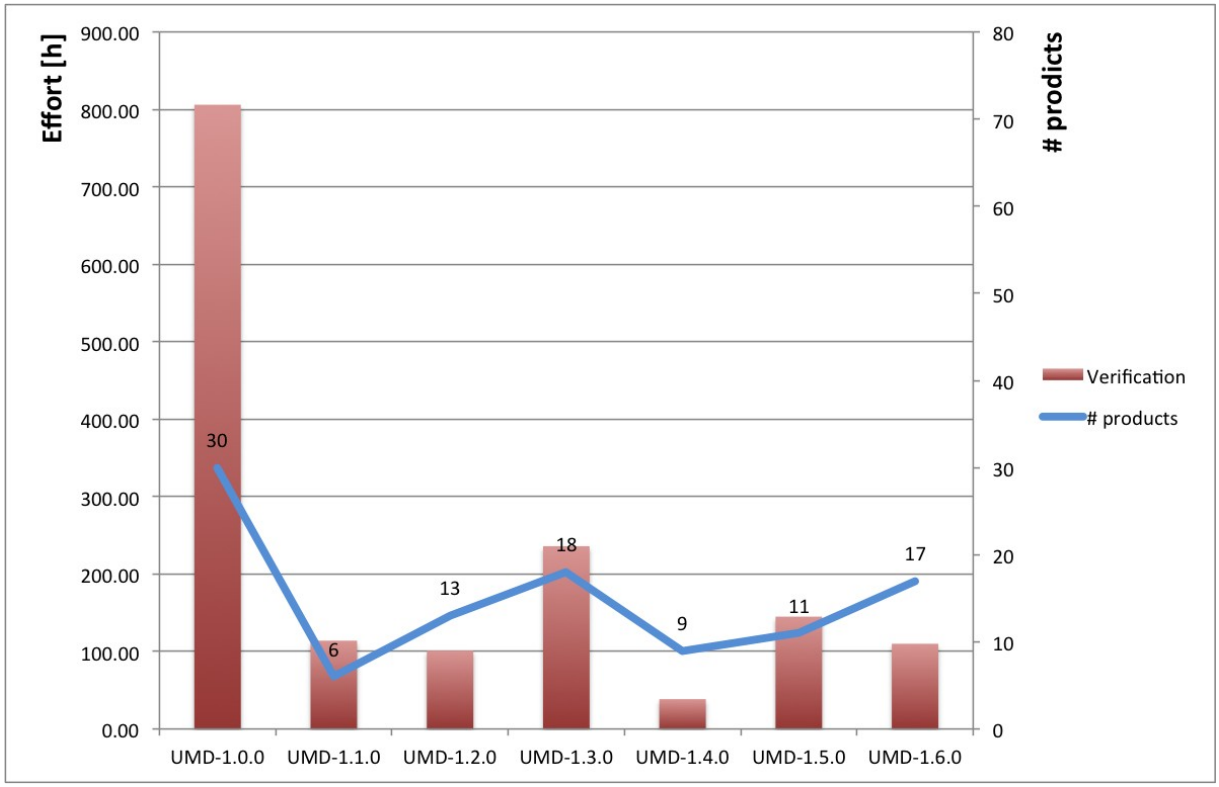

Figure 1: Seven UMD releases were published in PY2

Most of the products were verified within one working day[90]. However. There are considerable deviations in the median verification times reflect the very diverse complexity of software products that are provisioned for the production infrastructure.

2.2.3.3 Deployed Middleware Support Unit

The DMSU work was well established in PY1 and has evolved gradually during PY2. At the beginning of PY2 the leadership of the task was formally taken over by CESNET. Updated processes of DMSU were developed[91]. The changes address the most serious issue, the low ratio of tickets solved in DMSU which improved in PQ7 & PQ8 as procedures were reviewed following the experience gained in PQ5 & PQ6. The processes make sure that the remaining tickets, which are reassigned to the 3rd line support, are real defects requiring the Technology Provider taking action to provide a fix.

Further updates on DMSU work are related to the follow-up of tickets that were handed over to the 3 rd line middleware support (the Technology Providers). The discussion and an initial proposal were

[90] https://documents.egi.eu/document/1015

[91]https://documents.egi.eu/documents/504

presented at the 7th TCB meeting[92], followed by a convergence on an agreed process on how the DMSU will direct and manage the workload of all Technology Provider 3rd level support units, currently described in the DMSU wiki[93]. The upcoming MS511 will formally document the state of that process, which will be maintained in the EGI wiki on a daily basis.

The number of tickets in the table below reflects the work of DMSU:

Metric	PQ5	PQ6	PQ7	PQ8
Assigned to DMSU	218	183	173	212
Solved by DMSU	37	27	53	53
Reassigned to TPM	20	18	23	22
Assigned to 3rd level support	175	137	118	139
Mean (median) time to solve in days	24 (4.2)	17(11)	21.4(5.1)	11.3(1.9)

Table : DMSU GGUS ticket statistics in PY2

The numbers provided in Error: Reference source not found show that the incoming traffic reached a more or less sustained level. Error: Reference source not found shows the monthly distribution of tickets assigned to the DMSU. The apparent peak in June is related to the pre-release testing of UMD 1.0.0, which was routed through DMSU as well, the peak in February 2012 corresponds to adding WMS, a large complex piece of software coming in major upgrade compared to previous distributions, into UMD. On the other hand, there is a quiet time at the end of year.

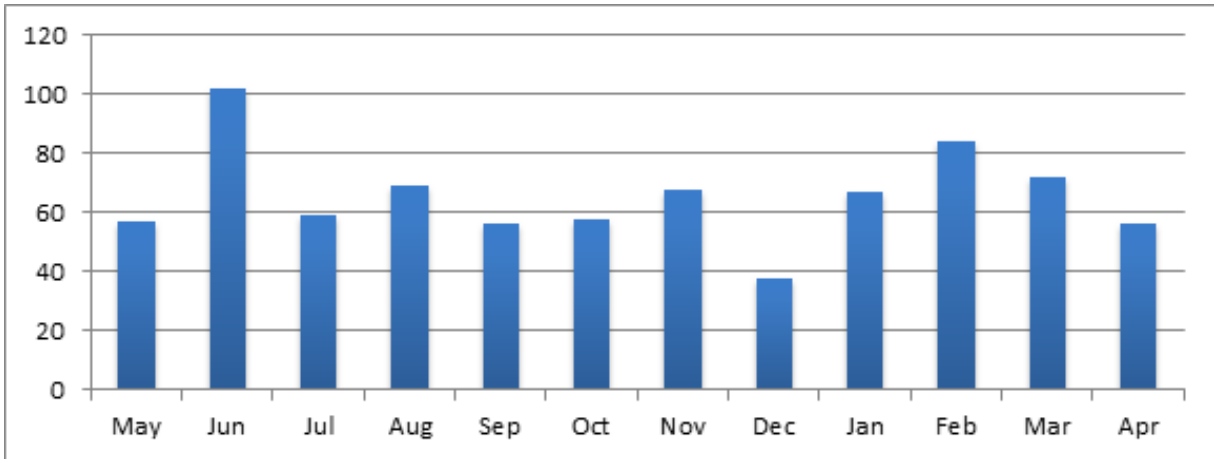

Figure : Monthly breakdown of tickets assigned to the DMSU for PY2

The tickets that are assigned back to TPM are, in more than 50% fraction, related to batch system support units. This software is supported internally in EGI; therefore the tickets must be routed through TPM instead of assigning to the 3rd line support units directly. The remaining issues assigned to TPM (less than 5% of all DMSU traffic) are operational problems wrongly considered to be

[92] http://go.egi.eu/TCB-7

[93] http://wiki.egi.eu/wiki/EGI_DMSU

middleware defects. This is an acceptable error rate at TPM. The relatively high average time to solve is caused by a few outlier tickets, being kept on hold until another problem is fixed. Therefore the median time is more representative.

DMSU regularly interacts with EGI Operations. The most common communication channel is the EGI Helpdesk, where issues get reported, specific issues are discussed through email or occasional phone calls between the DMSU leader and SA1 management. Vice versa, DMSU started to maintain the wiki page https://wiki.egi.eu/wiki/Middleware_issues_and_solutions where issues, which potentially affect wider user community or more sites, are thoroughly documented. In addition, the gathered DMSU expertise was leveraged in producing "best practices" manuals on BDII, WMS, and VOMS as requested by the Operations community. These documents are embedded in the EGI Operations Manuals wiki[94].

2.2.3.4 Support Infrastructure

During PY2 SA2.4 finalised the development of the 2nd iteration of the EGI Software Provisioning Workflow, which is now running in production mode since July 2011. The system has proven itself by supporting 7 UMD Releases, 6 CA trust anchor updates and 11 SAM Updates. Anticipating future developments in EGI and specifically in SA2, independent dedicated instances of the StratusLab marketplace and the StratusLab Appliance repository[95] were deployed and made available for SA2 and the EGI Federated Clouds Task Force to test the acceptance and uptake in the EGI community, before committing significant resources for its maintenance.

Request Tracker (RT)

During PY2 the software provisioning workflow was implemented in RT as the central tool coordinating ancillary actions and repository updates. This includes integrating the RT queue with the EGI Helpdesk for Technology Providers to make their recurring software releases available for provisioning, issuing digests for Technology Providers to monitor the progress of the provisioning activity within EGI. RT also collects a set of metrics related to the provisioning efforts for software and related to the software quality.

Repository frontend

During PY2 effort was put in to improving the mechanism that allows for automated publication of internal products releases into the repository.egi.eu portal. In addition, this mechanism was extended to include UMD Candidate Releases and UMD Releases. Another priority was offering general support for the repository.egi.eu Wordpress site including administration tasks such as software upgrades, support for the site's replication mechanism.

Repository backend

The repository backend is the subsystem of the EGI software repository that handles the business layer of the New Software Release workflow. It is responsible for sanitising the input coming from RT, for storing new software releases in its data store so that it can be passed through the workflow according

[94] https://wiki.egi.eu/wiki/Operations_Manuals

[95] http://marketplace.egi.eu and http://appliance-repo.egi.eu, respectively

to instructions given by the RT and finally for constructing the UMD releases based on the products submitted by the technology providers.

A new tool, the UMD Composer, was developed to aid in the management and publication of UMD releases in the EGI Software repository (more information on the Composer is given in D5.6). This tool has been used to produce all UMD releases in PY2. It was later extended to allow the publication of release candidates before the production repository is updated, allowing participants in the software provisioning process to do last installation checks. In preparation for extended OS platform support in EMI, and by request of the EGI operations community, support for SL6 was added, re-using the available infrastructure put in place for SL5 support. To support Debian 6 (as requested by the Operations community), a second repository management system had to be deployed next to the Red Hat-based SL5 and SL6 infrastructure.

IT Support

IT support faced mainly routine user support, e.g. setting up email lists and groups, changing configuration of the various systems (MediaWiki, DocDB, OpenCMS) to support new features, updating inspire-members from an Excel file send by e-mail, deleting or merging SSO accounts, creating new web sites for the EGITF an EGICF. The following list provides an overview of the most important tasks carried out in the second project year:

- Support X509 certificates authentication in Mailman and Wiki.
- Mirror the EGI wiki to ASGC in Taiwan.
- Create the "Early adopters" web application to manage the groups of TSA1.3 Early Adopters. The groups are then used in the Staged Rollout phase of the software release workflow for some automation. The application is available at https://www.egi.eu/earlyAdopters/
- Implement a news section in the EGI web site.
- Major change in the setting of Google Analytics to better track web visitors of all EGI web sites.
- Added links to EGI blog items on the EGI main web page.
- Provided CSS with printing layout for the EGI web pages.

2.2.4 External Relations and Community Engagement

During PY2 the original External Relations (NA2) and User Community Coordination (NA3) work packages were merged into a single new NA2 work package renamed Community Engagement. The revised focus on Community Engagement would be carried out through tasks covering communications and marketing, strategic planning and policy support, community outreach and technical outreach to new communities. These tasks built on top of work that had been primarily carried out in the old NA2 and included elements from NA3 that did not include NGI based User Support. This responsibility was transferred from the old NA3 to SA1 alongside the approval and management of VO regis-

trations and the operation of technical services for VOs (primarily VO-specific monitoring infrastructures).

The Marketing and Communications team (see Section 3.2.4.1) and the Strategic Planning and Policy Support team (see Section 3.2.4.2) that were primarily based at EGI.eu were given a revised remit and configured to support activities taking place in the NGI. The User Community Support Team operated as a single unit until the end of PM18, when it was restructured into two subgroups:

- 'Community Outreach' team is focused on community outreach through organising and participating at community events (see in Section 3.2.4.3).

- 'Technical Outreach to New Community' (TONC) provides community outreach through the development, support and coordination of technical services for new communities, and reported about in this section. This group – through its partners in Greece, Portugal, Spain and the UK is also responsible for the further development and provisioning of Technical Services for user support and outreach.

The NGI effort in the merged NA2 and NA3 was to be applied to the objectives of the new NA2 work packages through two new structures:

- NGI International Liaison (NILs): Individuals within an NGI given the responsibility for liaison on community agreed activities between EGI.eu and other NGIs, and the experts within the NGI.

- Virtual Teams: User to provide a flexible mechanism for bring experts from the EGI.eu central teams and from staff within the NGIs (assigned through the NILs) to work for 3-6 months on community agreed

These structures were presented to the new network of NILs at a workshop in November 2011 in Amsterdam, together with suggestions for initial projects. The process for proposing and setting up a Virtual Team is outlined on the wiki site and this was revised following a meeting at the EGICF12 which provided feedback on the new NA2 structures. The network of NGI International Liaisons is now well established, and all communication with NGIs around non-operational issues is routed through this network. There is an active core of around 26 NGIs who participate regularly in Virtual Teams. Effort is reported into the Virtual Teams through TNA2.1N if the staff members are NILs or Deputy NILs, and through TNA2.6 if they are not NILs but instead members of the NGI. The TNA2.6 task corresponds to the NGI Competency Centre. Staff members in NGIs reporting effort within the Virtual Teams, who are not NILs or Deputy NILs are therefore contributing to the NGI Competency Centre within the project.

Currently (March 2012), there are 2 Virtual Teams that have concluded, 9 teams in progress, and a further 7 proposed Virtual Teams in the process of being set up. More details are in D2.14[96].

[96] https://documents.egi.eu/document/1069

2.2.4.1 Marketing & Communication

Main website and wiki

The new website for EGI featuring new images, a new structure and enhanced social media interactivity was launched in March at the EGICF. The redesign process has included creating a new look and feel, introducing more visual elements and giving a greater prominence to social media channels. The new content structure is based on a dynamic two tier menu, which focuses content into strands for new users, existing users, the general public, policy makers and others.

A number of use cases have been developed and made available on the website and published through other channels, such as blogs and articles. A case study on hunting viruses using the grid was published in iSGTW[97] in January 2012, and is currently one of the top rated stories on the iSGTW website. The Glossary, Staff pages, EGI-InSPIRE's Deliverables and Milestones and other pages have been updated as required throughout the year.

From the 1 May 2011 to 30 April 2012, over 82,000 people visited the site, corresponding to 199,400 visits and 814,000 page views. Since the start of the project, there have been over 100,000 unique visitors to the website, with more than a million page views. During PY2, around 40% were new visitors, and 60% were returning visitors. The "stickiness" of the website was quite high in that on average people viewed 4.1 pages per visit and spent an average of 6:30 minutes on the site. The bounce rate was around 51%. The peak in visits was seen during the week of 18-24 September, the week of the EGITF in Lyon, with 7500 visits during that week. The most popular pages were the homepage, the timetables for the main events and the conference websites, followed by the documents pages and the intranet page, which contains details of how to access EGI's management tools. The highest levels of traffic were from search engine traffic at 39%, referrals at 35% and 25% direct traffic. Most people who found the site were searching for EGI or one of the events.

The wiki site has continued to provide the internal engine for the project, and was restructured in November 2011 to reflect the new project structure.

Materials, press and publications

Four issues of the EGI Inspired newsletter were prepared during PY2. The project team also produced Directors letters on a monthly basis, which were distributed to the whole consortium. Articles about EGI were also published in the e-IRG Newsletter, *The Parliament* magazine, the IN2P3 Newsletter, the CLARIN newsletter, *Public Service Review: European Science & Technology*, *Supercomputing Online*, *HPC in the Cloud*, *International Innovation* and *Public Service Review: European Union*. Extended articles about EGI were published in the inaugural issue of *PanEuropeanNetworks: Science & Technology* and 15 EGI-related items were published in iSGTW, which in December and January corresponded to an average of nearly one article per issue. Some of these items published in iSGTW have been picked up by the general media, such as *Symmetry*[98], *Discover News*[99] and *Wired*[100] and commu-

[97] http://www.isgtw.org/feature/hunting-viruses-finding-needle-haystack

[98] http://www.symmetrymagazine.org/breaking/2012/04/18/listening-for-the-sound-of-science

[99] http://news.discovery.com/space/making-music-with-the-sounds-of-symmetry-120419.html

[100] http://www.wired.co.uk/news/archive/2012-04/19/particle-physics-music

nicated to their readership through articles or Twitter. (The Twitter accounts which follow iSGTW have a combined following themselves of 1.6 million people).
In PQ7, the communications team partnered with Public Service Review to produce a dedicated 8 page booklet about EGI, which included an article about the Digital Agenda Commissioner, Neelie Kroes by the Public Service Review: European Union Editor, a 4 page article on EGI and a full page advert for the Community Forum. This was distributed to 140,000 policy subscribers by email and in printed form at the SciTech Europe event in Brussels in November, plus at other events attended by EGI. A dedicated weblink from the newly re-launched PanEuropean Networks website[101] points to the EGI website for 12 months from January 2012, and web traffic from this source will monitored. The DVD of the SciTech Europe masterclass was also issued in PQ8 on the cover of *PSR: Science & Techology*, distributed to 140,000 readers across Europe.

Social media and videos

The team has also been working closely with the EGI-InSPIRE work package leaders to encourage regular contributions to the EGI blog from across the project and wider community. By the end of April 2012, 108 blog posts have been contributed in total. The EGI team worked with CESNET to publish the feed from the most recent blog posts on the home page of the website in time for the EGITF 2011.
Members of the dissemination team have also blogged for the GridCast[102] blog at events such the Innovation Forum, Brussels, the eScience2011 event in Sweden and the SARA 40th anniversary event as well as a number of EC policy events, such as Research in Future Cloud Computing[103] in April. A new social media strategy was developed in PQ7, creating dedicated social media feeds[104] for general, user community, tech/ops and policy through Twitter, Facebook and the blog. To publicise the revitalised Facebook page, a mascot competition was launched on 11 January 2012 to choose a mascot for EGI from designs and concepts submitted by the EGI community. A People's Choice and a Grand Prize winner announced via social media.
EGI also worked with a film production company, Een van de Jongens[105] to produce a series of videos called "Stories from the grid". These short, YouTube friendly videos aim to introduce some of the work being carried out using the grid. The first video "Episode 1: The cone snail and the hunt for more power anaesthetics[106]" was released on 29 February 2012 at the ISGC2012 meeting to accompany the WeNMR keynote by Alexandre Bonvin and has 1000 views on YouTube. Two further videos in the

[101] http://www.paneuropeannetworks.com/

[102] www.gridcast.org

[103] http://gridtalk-project.blogspot.com/2012/05/cloud-computing-at-cutting-edge-of.html

[104] http://www.egi.eu/news-and-media/press/social_media.html

[105] http://www.eenvandejongens.nl/

[106] http://go.egi.eu/conco

series are in preparation on high energy physics and the lost sounds orchestra to be launched in May and June 2012.

Events and marketing to new users

The EGI.eu communications team worked with local NGI contacts to present the project at events throughout the year. These included booths at ISC2011 in Hamburg, Germany in June, in collaboration with KIT and at HPCS 2011 in Istanbul, Turkey in collaboration with the team at TUBITAK. EGI also hosted a booth at the HealthGrid 2011 event in Bristol, UK, run by the User Community Support Team and hosted a joint booth with e-ScienceTalk at eChallenges in Florence in October 2011.

The communications team added an EGI postcard to the delegate pack for the first European Gender Summit in November in Brussels, and participated at the event, which led to EGI signing the European Gender Summit Policy Manifesto[107]. Activities also included hosting a booth at SuperComputing'11, Seattle, US in November, an event which attracted around 10,000 delegates. EGI distributed T-shirts which featured a QR code leading to a hidden "Easter Egg" page on the website. EGI also attended and blogged from CloudScape IV in Brussels in February and hosted a booth, participated in the Programme Committee and presented at ISGC2012 in Taiwan, as well as filming material that appeared on the 10 year anniversary DVD. The team attended and blogged from the Girls, Women and ITC event in Amsterdam, and hosted information booths at both ICRI2012 in Copenhagen and the EGU General Assembly in Vienna in April. The communications team produced an earth sciences brochure for the EGU event, which included case studies in this area, and was distributed to booth visitors among the 11,000 delegates attending the event.

EGI Technical Forum 2011

The EGITF 2011 was held from 19-23 September 2011 in Lyon, France. The event was organised in collaboration with local hosts CC-IN2P3 and France Grilles. In total, 655 participants registered for the event, the best attended EGI event to date. There were 132 contributions from 296 speakers and 34 session conveners. The event was collocated with the Open Grid Forum, Grid2011, GlobusEUROPE, a French Grid Day and the 9th e-Infrastructure Concertation meeting, which was organised by the e-ScienceTalk project in collaboration with the EC. The US *HPC in the cloud* Editor attended the conference and posted several articles in HPC in the Cloud based on interviews at the event, including videos.

During the event, there were 250 Tweets from 27 people, 20 photos on Flickr, 27 blog posts on GridCast, including 9 videos. Before the event, the website received 3000 unique visitors, spending nearly 3 minutes on the pages, with 30,000 pages viewed. Nearly 1000 of these visitors went to the website during the event and 25% of these were new visits.

A feedback survey was launched at the end of the event and received 114 responses. Nearly 90% of respondents found the conference website quite or very useful. Around 75% found the registration process quite or very easy to use and around the same percentage found the EGI staff helpful. For the social media channels, around a quarter referred to the EGI blog, 14% to iSGTW, 43% read the GridCast blog, 19% looked at the Facebook group and 45% used Twitter.

EGI Community Forum 2012

[107] http://gender-summit.eu/index.php?option=com_content&view=article&id=278&Itemid=42

The EGICF 2012, was held in Munich on 26-30th March 2012, held in conjunction with the 2nd EMI Technical Conference and co-located with the 2nd Annual European Globus Community Forum. The event was advertised in *HPCwire*, *HPC in the cloud* and *Datanami*. Journalists attended from iSGTW and the US, leading to articles on GridCast, "Cyber-Bridging[108]", "UK-NGS Campus Champion – US/XSEDE Ripple Effect[109]", "Dealing with the Data Deluge: A silver lining is in the Cloud[110]" and "What are computers for?[111]". The event was themed around five tracks: Users and communities, Software services for users and communities, Middleware services, Operational services and infrastructure, and Coordination and Communication.

The event was attended by 518 participants and featured 171 contributions from 208 speakers. The marketing team produced a combined printed programme and book of abstracts. During the event, there were over 400 Tweets from more than 60 users, including the EGI news, policy, users and technical twitter feeds. Around 220 photos were uploaded to Flickr using the egicf tag and over 150 delegates downloaded the Conference4me conference app. The GridCast blog brought together 10 bloggers, posting more than 30 posts and 14 videos from the event. The event website received over 5380 visits from 1700 visitors. A Zoomerang survey was issued to the delegates, and 122 responses were received. In total, 93% found the conference website useful or very useful and 96% reported that registration was easy to use. 90% reported that the conference staff were helpful, or very helpful. For social media channels, of those that responded, 52% used Twitter, 25% Flickr, 18% viewed YouTube, 30% the GridCast blog and 45% the EGI blog.

2.2.4.2 Strategic Planning and Policy Support

The DoW changes introduced in PY2 expanded this task from "Policy Development" to "Strategic Planning and Policy Support". The scope of the task has been expanded to include activities in the area of strategy planning to better support EGI's strategic decision-making process.

Strategic Planning & Sustainability

The Strategy and Policy Team (SPT) has supported the development of the EGI Strategic Plan through a number of key contributions such as a market segment analysis, a detailed value network analysis, a EU2020 impact analysis, a revision of the EGI Ecosystem description and related SWOT analysis for the different roles. The SPT also supported the definition of a strategic measurement framework[112] for EGI based on the balanced scorecard[113]. The SPT supported the definition of business models through the lean canvas tool for EGI which included the concept of business space with exemplary business

[108] http://gridtalk-project.blogspot.com/2012/03/cyberbridgesmagically-built-from-middle.html

[109] http://gridtalk-project.blogspot.com/2012/03/us-xsede-ripple-effect-crosses-pond-uk.html

[110] http://gridtalk-project.blogspot.com/2012/04/dealing-with-data-deluge-silver-lining.html

[111] http://gridtalk-project.blogspot.com/2012/04/what-are-computers-for.html

[112]https://documents.egi.eu/document/1071

[113] http://en.wikipedia.org/wiki/Balanced_scorecard

models for EGI.eu, technology and resource providers[114]. This work cumulated in the EGI Sustainability plan[115]. Part of this plan may include the ERIC legal framework, which following the analysis made during PY1, the SPT prepared and submitted a number of questions to the ERIC Team for clarifying some aspects relevant to EGI. The ERIC team provided general answers to these questions[116].

Contribution to EU2020 and other EU Policy Activities

During PY2, the SPT has updated the list of EU2020 strategy actions relevant to EGI and analysed the individual contributions that EGI can make in these areas. The EGI impact on EU2020 has been derived in the area of the Digital Agenda for Europe[117] and Innovation Union[118] flagship initiatives. During PY3, the indicators will be improved in order to create a solid measurement framework to integrate with the EGI Balanced Scorecard.

On the EU policy aspect, the SPT provided continuous feedback through the participation in a number of surveys, either by answering the online questionnaires or by submitting position papers (e.g., survey on cloud computing, Common Strategic Framework for Research and Innovation, European Research Area).

In the area of socio-economic impact assessment, the SPT participated in the study "Development of Impact Measures for e-Infrastructures"[119] by contributing to the survey and participating in the dissemination workshop, in the Osiris project study[120], in the e.nventory data collection[121] and also in ERINA+ impact assessment through accessing their beta-testing phase for the impact assessment tool[122].

Communication

[114]https://documents.egi.eu/document/1040

[115]https://documents.egi.eu/document/1147

[116] https://documents.egi.eu/document/1011

[117] https://wiki.egi.eu/wiki/Europe_2020_actions/DAE

[118] https://wiki.egi.eu/wiki/Europe_2020_actions/IU

[119] http://www.ri-impact.eu/

[120] www.osiris-online.eu/

[121] http://www.enventory.eu/

[122] http://www.erinaplus.eu/

The SPT has regularly produced policy-oriented messages through several channels: eight articles on the EGI blog[123], one article in the iSGTW newsletter, three articles in the e-IRG newsletters and regular contribution to all editions of the EGI Inspired newsletter. Furthermore, the SPT has regularly provided updates through social media such as Twitter and Facebook.

The website and wiki of EGI have been regularly updated to reflect the SPT's activities evolution (e.g., by creating dedicated collaboration pages for each new collaborating partner, by creating a section for 'EGI in Europe', by revising the description of strategy and policy team, by refactoring and updating the SPT wiki section).

EGI Compendium

The Virtual Team project "EGI Compendium"[124] was created 5th December 2011 to define the set of strategic information to structure a body of knowledge for describing NGIs/EIROs in order to increase the transparency about their organisational structure, status and plan; to support strategic planning and to provide more clarity about the EGI.eu members activities outside the EGI community. The Virtual Team was led by the EGI.eu Strategy and Policy Manager and had active representatives from the Irish and Moldovan NGIs. The outcome of this activity was a questionnaire to map aspects such as identity and contact, strategy, governance, funding of NGIs, funding and staffing of the NGI/EIRO infrastructure, policy, outreach, services, users, infrastructure and technology. A first iteration of the data collection was executed and the first edition of the EGI compendium will be produced early in PY3.

Contribution to EGI Events

The SPT supported the organisation and running of several events through various functions such as chairing sessions, note taking, presenting and post-event reporting. The target events were the User Virtualisation Workshop, the first NGI International Liaisons Kick-off Meeting and the User and General EGI Sustainability Workshop. The SPT organised and chaired policy related sessions during these meetings:

- EGTF 2011: the EGI Policy workshop focused on disseminating the latest developments in the area of Horizon 2020 and Structural Funds through two hand-outs[125], presentations and discussions; the EGI Sustainability and Business Models workshop discussed the sustainability aspects of the EGI ecosystem entities (user communities, technology providers, NGIs) and an analysis document based on a survey was given to the attendants while a summary report was written afterwards[126]

[123] http://www.egi.eu/blog/categories/policy/

[124] https://wiki.egi.eu/wiki/VT_EGI_Compendium

[125] https://documents.egi.eu/document/799 and https://documents.egi.eu/document/800

[126] http://go.egi.eu/egitf11-sustainability-workshop

- EGICF2012: Sustaining the EGI ecosystem workshop covered various initiatives contributing towards a stable EGI ecosystem including presentations and discussions regarding the EGI2020 strategy and evolution of the EGI business model.

Supporting the Policy Groups

The SPT has provided secretarial support to the various EGI policy groups[127]. Some terms of reference (ToR) was revised to include the possibility for groups to appoint a task force (e.g., based on this, the TCB has created the EGI Federated Clouds Task Force[128]). The number of meetings held in PY2 varied from group to group, but in general the activities intensified. In fact, the number of meetings increased from 27 in the PY1 to 38 in PY2. Since January 2011, all policies and procedures created within EGI by any of the EGI Policy Groups have been approved under the agreed Policy Development Process (PDP). During PY2, 2 policies and 6 procedures were approved. In addition, the SPT has provided a draft escalation section for the Operation Level Agreements (OLAs) and aligned appropriate copyright statements.

EGI Glossary

The EGI common glossary has been updated to reflect the latest concepts emerged through the strategic activities and user community support. This will be marked as final during PY3.

External Collaborations

The negotiation and execution of memoranda of understandings (MoUs) intensified during the second year with the establishment of 16 new MoUs: CHAIN (Project), DECIDE (Project), e-ScienceTalk (Project), gSLM (Project), LSGC (VRC), MAPPER (Project), ScalaLife (Project), SIENA (Project), CSIR MERAKA (Resource Infrastructure Provider), HMRC (VRC), EDGI (Project), e.nventory (Project), SHIWA (Project), BCC/Ukraine (Resource Infrastructure Provider) the WLCG (VRC) and SCI-BUS (Project). Two more MoUs are under active negotiation: DANTE (Organisation) and UVACSE (Technology Provider). The progress of activities within the current MoUs were monitored through checking the defined work plan and the progress was recorded and accessible for the partners while all the reports have been stored in the EGI document repository.

Security

In the area of security, the major achievements were: approval of two new policies (Service Operations Security Policy and Security Policy for the Endorsement and Operations of Virtual Machine Images); successful "public information" session at EGITF; lead authorship for the production of a new Guidelines document addressing Attribute Authority Service Provider Operations; participation at OGF and TAGPMA meetings; leadership of the "Security for Collaborating Infrastructures" activity that resulted in Security policy framework draft between EGI, WLCG, OSG, DEISA/PRACE, and TeraGrid/XSEDE; participation in a new activity on Federated Identity Management for Scientific Collaborations; supporting the revision of the WLCG security risk assessment.

[127] http://www.egi.eu/policy/groups/

[128] https://wiki.egi.eu/wiki/Fedcloud-tf:FederatedCloudsTaskForce

Concerning IGTF and EUGridPMA activities, the biggest achievements in PY2 were: participation in the three EUGridPMA plenary meetings, three IGTF coordination meetings with the other continental PMAs in the Americas and the Asia-Pacific and in the CAOPS working group; participation in development of EUGridPMA policy and practice statement for technical controls on authorization services; an auditing process for the quality and integrity of the authentication operations is in place; attendance to the APGridPMA meeting and the EEF FIM workshop in Taipei, where agreement was obtained for the new AA Operations guidelines and changes to the authentication minimum requirements for Classic identity providers; participation in the IGTF and CAOPS Working group, where it was agreed to update the technical certificate profile to a OGF standard document "Proposed Recommendation". EGI representatives were the key driving force within the IGTF in identifying and coordinating requirements submitted to the EGI technology providers in order to identify software implementation obstacles related to new standards in the authentication domain.

2.2.4.3 Community Outreach

For the first half of PY2, community outreach activities were partly coordinated under NA3 but with some aspects, notably events, handled under NA2. For the latter half of PY2, community outreach was handled under NA2.4 with a focus on events. In the interests of continuity and comprehension, the narrative thread of the report is contained within one section of this periodic report.

During the first half of the year three new Virtual Research Communities were formally created through the signing of Memorandums of Understanding. These new VRCs were: Life Sciences Grid Community (LSGC); Hydro-Meteorology Research Community (HMRC); the Worldwide LHC Computing Grid (wLCG) and the Humanities represented by the CLARIN and DARIAH ESFRI projects in the form of a Letter of Intent (LoI). Other communities, namely: Astronomy, Computational Chemistry, Earth Sciences and Fusion continued to participate in the regular UCB meetings and other discussions but were not able to form themselves into VRCs. It is worth noting that LSGC and wLGC represent the two of the most successful models for VRCs. WLGC are a tightly formed centrally organised community of researchers spread over a wide geographical area but with many common needs and well established solutions being developed from within that community. LSGC on the other hand established their own original terms and processes specifically for their new organisation. The LSGC steering group meet regularly, circulate minutes to their community and select a spokesperson to act as the interface with EGI.

The purpose of the two forums, which are subtly different in focus, is to bring members of EGI's stakeholder community together in order to allow them to showcase their achievements and plans, as well as to stimulate future collaborations and partnerships. The focus of the EGICF is on the interface between research activities and the underlying infrastructure. By interface we refer to the scientific gateways, portals and workflows as well as their developers and managers, the main users of these gateways, and the providers of the services that lie behind these interfaces. The focus of the EGITF on the other hand, is on the technology that forms the infrastructure. In order to support these two approaches, the organisational processes that generate the programmes are different. The EGICF is assembled from proposals that have emerged from the community which have been submitted against the prescribed track descriptions. The EGITF is similar except that a more detailed programme

framework is prepared by the track leaders in order to ensure that the technological priorities identified by EGI are covered. This framework thus guides and informs the authors on their submissions. EGI also worked in collaboration with other European e-Infrastructure providers and their user communities to ensure that the meetings are infrastructure neutral in order to promote the harmonisation between European e-Infrastructures.

In addition to the bi-annual forums, EGI-InSPIRE also participated in and supported various other meetings organised by the user communities as well as technology-focused events targeted at research communities. These included:

- HealthGrid Conference, Bristol, UK, 27th July, 2011
- Astro meeting, Paris, 7th November, 2011 (VRC meeting co-located with Astronomical Data Analysis Software & Systems)
- Cracow Grid Workshop, Cracow, 7th – 9th November, 2011
- MAPPER school, London, 1st February, 2012
- SHIWA workshop, Budapest, 9th February, 2012
- RAMIRI, March, Prague, 12th – 14th March, 2012
- European Geoscience Union (EGU) General Assembly, Vienna, 23rd – 27th April, 2012
- HealthGrid May, 2012 (participation planned)
- International Workshop on Science Gateways (IWSG), Amsterdam, May, 2012 (participation planned)

Representatives from EGI-InSPIRE attended all of these events and contributed to the discussions and outcomes of the meetings. For HealthGrid, Cracow Grid workshop and EGU, Steve Brewer gave invited presentations which outlined EGI's role and offerings to these particular communities. The key follow on from this event will be for EGI.eu to continue to work with the Polish NGI to support their users in the wider context where applicable. The Astro planning meeting led to an increased coordination within the Astronomy community and was attended by Gergely Sipos. The goal of this event was to reorganise the Astronomy community to be able to proceed with EGI engagement, possibly in the form of multiple even more specialist VRCs and it was decided that a number of VRCs will be established. The MAPPER school in London was attended by Neasan O'Neil where the project is making good progress with their proof of concept scenarios, but still has some way to go before these solutions, which are indeed welcome by their user communities, can be properly bedded down in the production infrastructure.

RAMIRI stands for 'Realising and Managing International Research Infrastructures' and is a project funded by the European Commission under FP7. The RAMIRI project aims to deliver a training and networking programme for people involved in planning and managing international research infrastructures within the European Union (and Associated States). The focus of the project has been to organise two pairs of workshops designed as knowledge transfer opportunities for the target audience. Steven Newhouse attended one of the events of the first pair of workshops and Steve Brewer attended the Prague workshop. The event was attended by a few individuals already known to EGI but

many more from emerging or planned Research Infrastructures who were not so familiar with EGI and its role. The discussion sessions provided a number of useful opportunities to incorporate EGI's role into the case studies and scenarios. This RAMIRI workshop also helped form a number of closer working relationships with a number of ESFRI projects including the EISCAT_3D project and EPOS in addition to promoting EGI's role to the RI coordinators and managers present. Another follow on from this meeting was an introduction and link up between the financial work packages from CTA and EGI-InSPIRE to share best practices.

For the European Geosciences Union event, a Virtual Team was set up to coordinate EGI's contribution to the programme. The benefit of this was to draw in contributions to the event planning from partners as well as raise awareness of this activity in advance. In addition to a lively and well attended session on grid, HPC and cloud computing, the event also permitted the first formal meeting with the coordinator of the VERCE seismology project. An agreement was made to produce an MoU and align support activity on EGI's part with VERCE's planned deliverables. EISCAT_3D have also expressed a desire to discuss further how EGI could support their infrastructure that stretches from Northern Sweden to Oxford and beyond.

The other workshops and schools have also led to concrete follow on activity within the Technical Outreach team.

2.2.4.4 Technical Outreach to New Communities

The TONC group (established as a separate activity in PM19 to reach out and engage with new user communities, instead of supporting the use of existing user groups) worked in close collaboration with the community outreach team, as well as with other functional groups of EGI.eu.

2.2.4.4.1 Coordination

The activities and achievements performed/achieved by the TONC group of EGI.eu in PY2 were:

- Supporting the developers of the technical services in defining workplans based on feedback and requirements from user communities and their representatives. A 6 months-long plan was prepared for each service for 2011 and a second plan has been prepared for 2012 for AppDB and Training Marketplace. The 2011 work plans have been successfully implemented, and the services evolved as defined in the plans. There were minor delays in some of the items due to new, high priority requests that arrived from the community. Details are given in the technical service specific subsections below.
- Topical, technical workshops have been organised at the EGITF (Data Management; Portal/portlet/widget technologies) and in the Hungarian NGI (E-science workflows). These featured presentations from various NGIs, user communities, existing and possible technology providers, followed by discussions that clarified EGI.eu, NGI, VRC/VO/community roles in the further development and adoption of the technologies and influenced strategies in technical outreach.
- A dedicated section[129] on the EGI website has been setup to inform the community about EGI-related web gadgets and about knowledge that is available within community to support the devel-

[129] EGI web gadgets: http://www.egi.eu/user-support/gadgets

opment of additional gadgets. By the EGITF gadgets were available for AppDB, Training Marketplace, Requirements Tracker and these enable the integration of the centrally provided technical services into NGI, VRC, VO, institutional, personal or any other types of web pages. The GREIC service developer team from SA3 has developed two gadgets for their service in PY2 and this will be added to the EGI gadget portfolio in PQ9. Extensions to the AppDB and Training Marketplace gadgets started in early 2012 and will conclude later this year.

- The team supported the AppDB provider team to improve the quality of entries in the database, particularly (1) remove orphan entries whose owners/supporters no longer want to contribute, (2) merge entries that are about the same application but by different VOs, countries or support teams and (3) identify and update entries that include broken links in their references to publications, science gateways, download packages or any other online resources. Although the number of entries decreased as the result of the (1) and (2) actions, this decrease was lower than the number of additions in PY2, so the overall number of items stored in AppDB increased. (See details in 'Uptake' section of the Applications Database report below.)
- In collaboration with the technology providers in TCB, the processes around communicating and managing requirements between user communities and product teams were finalised. The process has been documented and a succinct summary of the whole requirement workflow has been made available alongside with open and solved user requirements through the EGI webpage[130]. During PY2, 31 user requirements[131] were resolved by this process. These include 3 MPI and 2 high-level topics that were flagged as priority by the User Community Board. 23 additional requirements[132] are still in the TCB process, at various stages:
  - The delivery of "documentation about WMS recommended configuration" (a topic considered as priority by the UCB in 2010) was already delayed twice, and it will not be provided as part of EMI-2 release. The date is not known when it will be provided.
  - 1 UCB topic and 1 "normal" requirement were endorsed by TCB and are waiting for technology providers to accept as item they want to implement.
  - 2 requirements (1 of them is a UCB topic) are still under evaluation by the TCB (since February 2012).
  - 2 requirements have been returned by TCB for clarification. These are followed up by UCST with the submitter communities.
  - 2 UCB topics and 2 'normal' requirements were assessed by technology providers and are waiting now for decision if they are going to be implemented or returned to the submitter.

[130] User requirements and process: http://www.egi.eu/user-support/getting_help/

[131] Solved user requirements: https://wiki.egi.eu/wiki/Solved_user_requirements

[132] Open user requirements at TCB: 3230 (UCB topic), 926, 1778 (UCB topic), 1780 (UCB topic), 2731 (UCB topic), 1626, 2733, 924, 910, 1777 (UCB topic), 3563, 3406,3404.

- 7 requirements have been identified as bugs through the EGI Helpdesk, Technology providers are currently assessing the cost of fixing these in future releases.
- 4 other requirement are still under investigation by the helpdesk support, to decide whether these relate to bugs, or missing features.

- Time was spent investigating and the seting up structured wiki pages about middleware services APIs; command line clients for non-Linux platforms; and MPI capabilities (API[133], CLI[134], MPI[135]). These extend and provide a more user-centric structure to the documentations that are provided by external technology providers or EGI-InSPIRE teams (e.g. MPI SA3).
- The EGI Blog[136] has been used extensively by the team to report about activities, achievements. 29 blog posts have been written by EGI.eu UCST members, and 2 by providers of Technical Services (used to be TNA3.4, now TNA2.5). While the team hoped to receive feedback from the community through these blog entries, this has not happened and the blog remained a one-directional communication channel towards the community.
- The team is involved in six, already active Virtual Team projects (Intelligence Collection; ESFRI Contact List; MPI; Federated Identity Providers Assessment, Fire and smoke simulation; Speech Processing on the Grid) and helped NGIs formulate additional VT templates (ESFRI Demonstrators; Top applications and their user communities, Application Porting Howto Guide)[137]. In case of enough interest from the NGIs, these projects will start in PY3.
- Together with the owners of "robot certificates" the TONC team established a process that could help EGI assess the size of user communities who access the infrastructure through applications and portals that use robot certificates, and therefore are invisible by the current statistics[138]. The full picture will be visible only in the next few months, however based on the initial findings these 50-55 robots currently serve around 200 users, which is still insignificant compared to the number of people who use personal certificates to access EGI. (around 21K)
- Collaborations with the ScalaLife[139] project cross-referenced user support web pages, training materials and on the identification of EGI users who could benefit from the application support that ScalaLife can provide for GROMACS and DALTON users.

[133] Table of EGI middleware APIs (not final at the time of writing): https://wiki.egi.eu/wiki/Service_APIs

[134] Command line user interfaces for EGI middleware services: https://wiki.egi.eu/wiki/User_Interfaces

[135] MPI User Guide for EGI users: https://wiki.egi.eu/wiki/MPI_User_Guide

[136] EGI Blog: http://www.egi.eu/blog

[137] EGI Virtual Team projects: https://wiki.egi.eu/wiki/Virtual_Team_Projects

[138] Statistics about EGI robot certificates and users: https://wiki.egi.eu/wiki/EGI_robot_certificate_users

[139] ScalaLife project: http://www.scalalife.eu/

- Established a "portal community" within EGI to provide a forum for those developers, operators and users who work or are interested in activities related to web based science gateway and portal frameworks/components. Through this community we collected information about existing science gateways and science gateway enabling technologies. The information has been registered in the EGI Applications Database and presented through gadget instances on a 'Science gateway section' on the EGI website[140]. The AppDB data and the website will evolve in PY3 through a new Virtual Team project, titled 'EGI Science gateway primer'[141].
- To simplify the development and management of complex applications on EGI, the team has collected information about workflow systems and workflows that exist within the community. Information about these have been registered in the EGI AppDB and is presented on a dedicated section of the EGI website[142].
- In late November an email list has been established for those EGI members, who are interested in and/or who already use the MATLAB software on NGI resources. The list has ~20 members and aimed to be used for information exchange between these partners, and as a discussion forum where the details of possible need for wider, more harmonised MATLAB access in EGI can be clarified. Such request has not emerged from the email list yet.
- Collaboration with SA1 and JRA1 on correcting the methods and tools by which the number of EGI users is counted took place. Particularly: those VOs that were inactive in the last 12 months have been suspended (so their users do not appear in statistics); Requirements in the further development of the Operations Portal have been captured for JRA1.
- Contributed to the development of the EGI Strategy through internal discussions in EGI.eu, as well as public discussions at the EGI Sustainability Workshop. The team's prime interest is in 3rd party integration & operation of Virtual Research Environments on EGI and supporting these by collecting and arranging software and technical support from the community for user communities.
- Joint work with SA1 has begun on defining a monitoring framework with documentations within the EGI service monitoring system that can be used by projects of the EGI ecosystem that provide user level software services. The development continues and expected to conclude in the forthcoming months, after the major release of the EGI monitoring system.
- Began to develop short, focused presentation about support services and activities that the EGI community provides for (new) communities. Share these presentations with the NGIs through the EGI Training Marketplace, with the goal that NGIs will use these when they talk to (or train) new communities.
- Work with EMBL-EBI from SA3 and CSC from the Finnish NGI on two application pilots that would integrate the 'European Genome-phenome Archive' service with federated identity management services from the EGI community.

[140] EGI science gateways: http://go.egi.eu/sciencegateways

[141] Science gateway primer Virtual Team project: https://wiki.egi.eu/wiki/VT_Science_Gateway_Primer

[142] EGI workflows and workflow systems: http://go.egi.eu/workflows

2.2.4.4.2 Applications Database

The activities during PY2 have been primarily focused on the following key areas: Architecture (revise codebase and RDBMS), Quality of information, Information retrieval, Notification & Dissemination, Cross-browser compatibility.

Architecture

More specifically one of the first tasks that the AppDB development team dealt with during PY2 was the migration of the AppDB DBMS system to PostgreSQL. The previous RDBMS of choice, namely MySQL with an InnoDB engine, suffered from limits with regards to scalability, performance, and query language expressiveness. On the other hand, a more advanced, yet open RDBMS such as PostgreSQL, now seems to be more suitable for the project's needs. Its superior query language expressiveness, paired with better performance and advanced features can deliver greater data transformation abilities inside the data tier. This automatically translates to better optimized code within the existing codebase architecture, and thus a better overall performance in general.

Service wide, the existing REST API has been extended with more resources, and now exposes everything that is needed in order to build a web site, or even an application based on it; there is no need for server-side access. Furthermore, the REST API has been also extended with authenticated access, in order to suppress sensitive data such as email addresses in anonymous calls, and to provide information about the authenticated user's permissions.

Quality of information

One of the main activities of the EGI Applications Database during the PY2 was to improve the quality of information stored within the database. This was achieved by adding a set of new functions to AppDB that support the community as well as the EGI.eu and AppDB teams in order to improve the accuracy of the stored information and to be able to maintain this high quality in the future:

Entry Problem Reporting System: There have been several cases, where a certain application entry raised disputes concerning the ownership of the referenced application, related logos or names. Users are therefore provided with the ability to report problematic entries with a short text about the type of infringement, so that appropriate actions can be taken to correct the problem. EGI.eu User Community Support Team and the owner(s) of the entry are notified via e-mail about such reports.

Broken Link Detection Notification System: A mechanism that is capable of detecting entries with broken URLs within their references section is part of AppDB. The system helps entry owners as well as the central content management team maintain the high quality and accuracy of the stored information. The broken link detection system scans the database periodically. In case a broken reference is found the system sends an email to the owner(s) of the entry as well as to EGI.eu. Extra mechanisms have been developed to monitor and remove invalid or non-existing references from application's extended information, such as publication links and external URLs.

Application Revocation Mechanism: In case a problematic entry is not or cannot be corrected, then users granted with a "Management" role within the system can revoke these applications from public visibility. Revoked entries do not appear in index pages, search results and are not included in future

statistics or export functionalities, etc. The entries still remain in the back-end database though and can be restored to their original state if needed.

User Comment/Ranking System: Quality of information can be a sensitive subject in large data stores. A module has been developed to help ensure content quality is the provision of quality in related metrics, such as comments and rankings. These metrics are provided by authenticated end users and are visible by everyone.

Merging duplicated/overlapping entries: Mechanisms have been developed for supporting the process of merging multiple instances (flavours) of the same application into a single one, thus avoiding confusion in the community about the authenticity and ownership of these applications. Applications can now be related to more than one discipline and sub-discipline, so that they may cover all of their possible uses. For maintaining the quality of information, each application's name is forced to be unique during registration and while editing an application.

Specifying people's relationship to entries: Scientific contacts can be tagged as experts on individual fields, such as a specific middleware or virtual organization, so that users may be led to the right person for further information.

Information Retrieval

As the number of entries stored in the system grows, finding one, or a set of applications or tools for a specific purpose becomes a non-trivial task. Despite the powerful filtering mechanism finding a search term that provides all the entries of the set is sometimes impossible. Tagging and tag based searches provide better mechanisms, allowing the community to create custom sets from the entries stored in the AppDB. Corresponding developments have been completed during PY2 of the project, thus providing a set of predefined tags attributed automatically by the system to certain entry types (e.g. in order to distinguish between native entries, and entries imported by other similar systems), as well as free-string tags which can be attributed by the users on demand.

Notification & Dissemination: A notification mechanism has been also integrated into the system. Users can now subscribe to mailing lists or RSS/Atom news feeds, so that they get informed about changes in AppDB's context. The mentioned mechanism is flexible enough to allow users to customize and constrain the amount and type of context they wish to be notified about.

Cross-browser compatibility: A significant effort has been made toward improving the user experience. The user interface has been also revised with a more content oriented approach, and information acquisition has been enhanced with a very powerful and flexible searching mechanism. The portal can now operate in all major web browsers (firefox, chrome, safari and IE9). Portal navigation is more direct and clear to help users focus on the group of applications they are searching for.

Dissemination tool

At the time of writing of this document, a new AppDB extension, the so called 'dissemination tool', is in its final review phase. This tool is intended to be used only by the management team and it main aim is to provide a means of collectively informing groups of AppDB users about important informa-

tion, such as community events, policy changes, required actions, etc.. The tool consists of three tabs. In the first tab, the manager using it can specify the set of criteria that the group of users to be informed must match. This can be done either through the use of what can be described as a graphical query builder – which is the default – or by directly entering the search criteria in text. The tool also allows for excluding individual persons from the result set as a means of fine-tuning the final recipient list. In the second tab, the manager may edit the message that is to be sent to the selected recipients. HTML and plain-text formats are supported, and the manager may choose to preview the message in the browser as well as externally before sending the final message to the list. Finally, the third tab provides a sortable, searchable log of all the sent dissemination messages, including such information as date, sender, recipient list, and text.

Community interaction

The AppDB has been extended with more community-based capabilities in order to motivate users to participate in information maintenance and to enhance the communication with the rest of the community in more direct ways. To this end, two new features have been developed. The first one is the ability of any registered user to request for his/her inclusion in an application's development/support team. The owner and the related contacts of the corresponding application get informed for such requests via e-mail and the application's database web portal in a centralized fashion and may accept or reject them accordingly. The second feature gives the ability to registered users to directly send an e-mail message to any related contact of a specific application. This way the user can ask questions about any application directly from the people involved with it.

Write enabled API

After much talk and requests from the community, the AppDB is moving towards providing write access though its REST API. Providing the community with write access will enable the provision of complete solutions from third parties, such as NGI-specific portals, will permit synchronization with similar databases from other projects, and may also provide the basis for write-enabled web gadgets. The portal itself is already relying exclusively on the REST API for read access, and once write access is complete, all of the portal's write actions will be exclusively served through the API as well. At the time of this writing, write access is still in alpha status, and is expected to achive beta status and be ready for review by July 2012.

Uptake

On the 14th of March 2012 the database includes the registration of 366 applications, 54 tools and 669 person profiles. This means an increase of 24 for applications, 18 for tools, 51 people in relative sense since May 2011. However, due to the merges (11 apps.) and removal (5 apps and 2 persons) of outdated entries, the real increase for applications and people is higher: 40 new applications and 41 persons got registered since May 2011. This is roughly 60% less new applications but 40% more tools

than got registered in PY1 (93 new applications, 10 new tools in PY1). During PY2 there were 21 releases[143] of the AppDB services, including one major, five minor and fifteen small updates.

Because of capability limitation, Google Analytics can monitor only the Web, but not the gadget instances of AppDB. Therefore a Piwik web analyser has been put in place in PQ6 for the AppDB service and this monitors both the central web instance as well as it's gadget deployments. Piwik statistics show a strong increase in visits (Q6→Q7→Q8: 3568 → 4888 → 6998), longer visitor sessions for both new and returning visitors (30→40→50 minutes for returning visits) and a stabilised bounce rate (~41%). The statistics include the traffic from the developer and content maintenance teams from IASA and EGI.eu, but yet they are promising signs of increasing uptake.

2.2.4.4.3 Training Marketplace

The Training Marketplace is a tool to view and advertise training events, resources and requirements. It is aimed at supporting cooperation between trainers and users in different localities and projects by connecting the groups through the activities that are established within the NGIs and scientific clusters. The goal is to enable users to achieve better scientific performance when using EGI and guide the establishment of self-sustainable user communities. Among the provided services include training events list, which allows trainers to advertise their training events and to be made aware of other training events being run within the community, a map of these training events, a repository of training materials and other resources and a web gadget that can be used to embed customised instances of these services into different websites.

During PY2 the Training Marketplace had several releases, some of these based on replacement of Drupal modules, resulting richer functionality and/or more intuitive interfaces to enhance the user experience.

Early in PQ5 version 2 of the EGI Training Marketplace was developed, tested and then released on June 7th. Version 2 of the Training Marketplace has a significantly new look and feel to version 1. It contained significant increased functionality which included a map view of NGI training events, a training wish list (a web form is available for users to specify their needs and publish a list view of these requirements), the ability to advertise expertise and ideas e.g., training resources (a web input form is available for people to advertise their expertise, ideas and resources and publish a list view of these) and a form to advertise relevant MSc, PhD and other taught courses. Version 2 also fully integrated approximately 9000 items of legacy material from the Digital Library training repository with Apache Solr as a search engine. As a consequence of this, there is now seamless search functionality across all objects in the Training Marketplace. For example a search on the term 'gLite' would retrieve both events and resources (e.g. materials, infrastructures, VOs) relating to the term.

For the remainder of PQ5 work continued on development of V3 of the Training Marketplace. The development process was rapid as known functionality was implemented into successive iterations whilst a reliable version was available as production service. The most significant development task for v3 was the creation of a web gadget (in fact three or four separate gadgets) that allow other projects to embed the Training Marketplace into their own websites and for the gadget to be

[143] http://appdb.egi.eu/?p=about:changelog

customisable. This customisation mainly applies to the browse gadget and it takes the effect of showing a default view of the local region. Another feature included in V3 is the ability to report inappropriate comments (through the user feedback channel) for the UCST to investigate and remove. Once a number of wishes/requests have been added to the database a tag cloud will be included.

All feature and functionality requests for the Training Marketplace are now managed through EGI's Requirement Tracker system. The Training Marketplace was regularly reviewed and new and refined requirements provided to the team in the UK. In order to promote the Training Marketplace it was decided that a brochure was not essential at this time as it is largely self-explanatory and intuitive to use. A bookmark was thought to be a more efficient promotional tool and a Training Marketplace bookmark has been prepared for the Technical Forum. The V3 major release was timed before the EGITF and involved the development and the release of the Training Marketplace gadget.

Another release in PQ6 provided admin interfaces to allow TM administrators (STFC and EGI.eu) to see events waiting for approval in a central place and to identify when admin need to login to view restricted parts of the website.

Since these releases development has focussed on integrating new Drupal modules into the TM website. These vastly improved the display of training events and allow greater modifications in future. Another new module has been investigated to capture user input when advertising, and a new resource type "Online resources" has been added to the development server ready for production in PQ7. Other significant improvements have been made to the user interface in PQ7 too. These include:

- Updates to the calendar display to support multiday events as single, clickable units.
- The calendar was made the default display option when browsing for events, enabled by using a new module.
- Bugs in the map integration (using the open-source Ammap tool) into the new module were fixed.
- When the cursor hovers over a country a pop-up displays past events as well as future events - five each by default and a 'More' button to view fuller content.
- More coherent and rationalised usage of 'Go Back' buttons across the different sections.
- A number of routine and security updates have been applied to Drupal modules.
- The Cookie Law has been extensively investigated for its impact on websites using Google Analytics (incl. www.EGI.eu.). Until a Drupal module is developed to enable full compliance the TM is adopting EGIs policy wrt Cookie Law.

In PQ8, a new version of the gadget has been released with improved functionality and flexibility for NGIs wishing to use it. New options have been released to allow the user to filter events by country, displaying only those relevant to their region. The events can be displayed on a calendar or in a list, customisable by colour, font etc to enable it to match the branding of the user's website. These are resizable to maximise their fit.

Other work has included removing the "Requirements" buttons from the Training Marketplace. This is because the function was not being used. The capture information regarding the number of training days offered through the TM, new fields were added to the events form. We have also been working on monitoring scripts that generate statistics on the number of users to the TM via the main interface in the EGi website, and via gadgets that have been embedded in other websites.

Uptake

Despite the continuous evolution of the Training Marketplace the number of unique visitors captured by Google Analytics on the site dropped significantly since May 2011. PQ5-425; PQ6-324; PQ7-197. The number of training days delivered by the NGIs has not dropped, so this trend cannot be explained with having fewer events.

The developer team is trying to reverse this trend in2012 by planning new functionalities for the Marketplace that could make the system and content stored within it more customisable and reusable for individual countries, projects, groups and use cases (e.g. through improved web gadgets; through tagging of stored items). Because Google Analytics is unable to capture the traffic through the gadgets, the negative trend in visitors of the central instance may be compensated by higher traffic on gadget instances. The team will need to put sufficient mechanisms in place to capture the traffic through gadgets. The open source Piwik service which is already used by the AppDB team can be such a tool.

2.2.4.4.4 VO Services

Until the end of PM18, the "VO Services" was part of TNA3.3, but since PM19 the same group (Ibergrid in Portugal and Spain) provides a Client Relationship Management (CRM) system for EGI from the TNA2.5 task. This section of the report provides information about the achievements of the VO Services task between May-October 2011.

The technical instantiation of a user community within the infrastructure is a VO. Setting up and operating a VO is a complex task that requires an important effort. Although various technical services are available to collect availability, accounting and monitoring information, as well as information that will improve operation of the VO, procedures are neither easily available nor complete. The EGI VO Services task aimed at supporting VOs in the whole process of start-up, management and operation, pointing out to tools, services, documentation and guidelines to maximize the usage of the resources and bridging the VO community with the infrastructure they need. The activities performed by the team between May-October 2011 are documented elsewhere[144]:

- Collaboration with SAM developers on providing a basic, Nagios-based, VO-specific and VO-centric testing and monitoring system for VRCs, and on extending this service with additional components and capabilities as the communities' needs evolve. The VO SAM is now a flavour officially supported by SAM, and will continue to be enhanced with the introduction of a new component, the Profile Management Database (POEM) substituting the Metric Description Database (MDD), allowing higher flexibility to describe existing metrics and group them in order to run dedicated tests. The team also operated VO SAM services for phys.vo.ibergrid.eu, life.vo.ibergrid.eu,

[144] https://wiki.egi.eu/wiki/Services_and_Tools_Portfolio

ict.vo.ibergrid.eu and fusion, and provided consultancy for We-NMR and HealthGrid (on behalf of the Life-Science Grid Community VRC).

- Development of specific tools that can address VO requirements focused on easy access to information and on the mitigation of data management issues namely the VO Admin Dashboard, a highly configurable integrator dashboard for project and VO tools, and the LFCBrowseSE, a tool to improve file catalogue views, currently in its third version, and extensively used by the LSGC VRC for SE decommission.
- Identify, evaluate and test third-party tools and services that emerging VOs can use or deploy. Document those services through white papers and manuals for VRCs who wish to operate such services themselves.
- Produce a summary of the relevant policies for VOs.
- Produce the first draft version for the removal of obsolete VOs from the infrastructure, including workflows and communication details among main actors.
- Produce a work plan for the VO operational portal including feedback from JRA1 and the Life Sciences Grid Community. The VO Operations portal is a tool foreseen to empower VOs with a service necessary to deliver the best possible availability for end users. The proposal is to make it accessible as a module of the EGI Operations Portal, with a central instance with an entry point per VO. The development for this portal has started on October 2011, and a first release of the tool is already available and under test by the communities which request it.

2.2.4.4.5 Client Relationship Management system (CRM)

The merger of the NA2-NA3 activities was driven by the need to give more focus and resources to outreach to new user communities. A coordinated outreach activity needs a tool where NGIs and EGI.eu can record contacts, leads, conversations. It was decided to setup a Client Relationship Management (CRM) system for this purpose, using the effort previously used for VO Services. This section of the report describes the achievements of the CRM task since November and its plans for PY3.

The first task for the team was to identify a CRM tool (among the several available open source sets) that could fit EGI use cases and context. Two different use cases were identified:

1.NGI International Liaisons should be able to record information about known projects involving their national institutes, record information about the individual contacts in each institution and track discussions held with those contacts. NGI staff must be able to cross-check experiences and observe what other NGI members did differently to bring on board a new user community.

2.EGI.eu members should be able to feed the system with new initiatives and contacts that they obtained through other top mechanisms. Simultaneous, EGI.eu members should be able to assess and track the progress made by the different NGI approaching their user community leads.

The team identified and installed[145] vTiger, a highly customisable well-known LAMP open source CRM technology at the Laboratório de Instrumentação e Física de Partículas. Lisbon, Portugal. The

[145] http://crm.egi.eu

details of the above two use cases were further elaborated by members of the "Intelligence Collection & Analysis" and the "ESFRI Contact list" Virtual Team projects and were used by the team to fine tune the configuration of the vTiger installation. One should note that given the commercial nature of CRM tools, this was not an easy job, and that some of the functionalities do not really map to EGI use cases. The main effort was delivered in the definition of the data models (structures and entities), in the implementation of new modules, in the customisation of the Web User interface and of the reporting module (creation of dashboards, creation of automatic reports, etc). Finally, the vtiger user authentication mechanism was tested and successfully integrated with the Single Sign-On (SSO) EGI authentication mechanism.
A pilot instance was presented and demonstrated to EGI.eu and to NILs on the 13th February 2012, followed by a period where users could use and test the system functionalities, and by some fine-tuning of the system. A number of NILs provided feedback during the test period to the providers. The final system was opened for the NILs and their partners during the Community Forum. The uptake of the service within the community will be known in PY3.
Since the EGICG some additional requirements were identified in regards with the traceability and auditing of changes of individual CRM records. According to the official vTiger documentation the new release of the tool may provide solution to some of these issues. An installation of this new release is currently under evaluation on a test site at Universitat Politècnica de València.

2.2.4.5 NGI Technical and Community Outreach

The following section provides information about the activities that were carried out by the NGIs outside of the Virtual Teams, but contributed to the outreach and support for new user communities.

Armenia: The Armenian NGI User Support Team besides the standard operation of user support services improved the existing support services and develops new tools and services to reach out and serve new communities. Regular meetings with ARMNGI potential user communities In meteorology, physics and life sciences have taken place and new scientific software packages (NAMD, SAGE, CROMACS, WRF) have been integrated to the computing sites.

Bulgaria: During PY2 the Bulgarian NGI leveraged the increased storage capacity to deploy and use more storage intensive applications from diverse scientific areas, especially environmental security and environmental protection. The high-performance Grid clusters that form the majority of the Bulgarian Grid resources facilitated the usage of advanced parallel applications that follow the MPI standard. In the last few months we added new nodes equipped with state-of-the-art NVIDIA M2090 cards for GPGPU processing and initial success has been achieved in porting applications from computational physics and computational chemistry achieving high power and cost efficiency. The results, achieved on the Grid, were presented at high-ranked scientific conferences, relevant to the respective communities, as well as on workshops that have been organised specifically to foster the collaboration of Grid users. Recently the Bulgarian NGI has had its 10th year anniversary meeting, where representatives from the member universities and research institutes discussed the next directions for expanding the user base and facilitate the access for scientists and Ph.D. students to advanced Grid resources. We expect that a concerted effort to improve the curricula of Grid and related subjects for master and Ph.D. students should contribute to bring on board new user

communities from domains like bioinformatics and national heritage. The current hardware resources of Bulgarian NGI are state-of-the-art and adequate for the needs of our current Bulgarian and international users. Our focus for the next period is to improve the utilization of the facilities and to promote the take-up of new technologies like GPGPU, cloud, etc. as part of our Grid applications.

Croatia: Since October 2011 the team started publishing monthly reports to keep the national grid community updated with status, activities and usage of the national grid infrastructure. Reports are written in Croatian language and contain important operations actions and usage statistics per institute in a given period. In December 2011 CRO NGI organized the annual meeting of the Croatian grid community – The Third CRO NGI Day. At the meeting grid community members presented their use cases, experiences and challenges with grid technology.

Cyprus: During PY2 CyGRID user support team had worked with users from the Statistical and Mathematical community that were interested in running the R project for statistical computing on grid. R is a language and environment for statistical computing and graphics. We have installed R in two Cyprian clusters: CY-01-KIMON and CY-03-INTERCOLLEGE. A number of dedicated training sessions took place for these users. CYGRID, has continued user support for the current running applications R project, SimpleScalar Simulator, Scheduler Sim, and NS2 for new users from the local scientific community. Standard support for Alice, Atlas, LHCb and Biomed VOs continued. Users from the Cyprus Institute are interested in using the following applications: MPI and OpenMP (LQCD calculations), LTools, UCNS3D (CFD Software), Metis/ParaMetis/ParMGridGen on both EGI and PRACE. The team aims to address this requirement in PY3.

Czech Republic: The user support team supported the VOs: ATLAS, ALICE, AUGER, BELLE, VOCE, MPI, MetaCentrum and disseminated their experiences to potential new users. VO Auger, the biggest non LHC grid CPU time consumer, continues to use CESNET servers for central services. The high amount of simulated data requires constant maintenance but also attracts more users. Recently we demonstrated usage of software for electronic structure calculations of solids on the VO VOCE infrastructure. Individual users and small team outside big international communities are also supported. After several months of negotiations with the Czech local ELIXIR site (institute of Organic Chemistry and Biochemistry of the ASCR, Prague) there was agreed that CESNET would support the local ELIXIR's site integration to the whole ELIXIR infrastructure by providing a starting computing and necessary network equipment. A MoU was signed between CESNET and Institute of Organic Chemistry and Biochemistry of the ASCRin November 2011. In March has formed a National Council with aim to build an official ELIXIR node in the Czech Republic, with participation of other institutions (CESNET, CEITEC, CERIT-SC, IT4I,...). NGI user support team also gained a new user community from RECETOX centre (integrated to the large infrastructure project CETOCOHEN), their special application software, (WRF - Weather Research & Forecasting Model) has been installed and tested for Grid usage. Concerning outreach to new users in the future: the team will organise a 3 day long Auger workshop in June and a Rexetox workshop in April or May 2012, and CESNET conference in April. Additional workshops for new user communities may be organised. Participation on the proposal for a new grid data management system for the VO AUGER continues.

Denmark: We have started a project to approach the Danish geophysical science environment in order to clarify which users and subjects are most likely to be potential users of GRID resources. Institutes

we have started to contact include mainly heavy satellite data users e.g. geodesy, geomagnetism, and oceanography. A secondary approach will target more broadly and also include other aspects of 'geophysics' such as e.g. life-sciences. Our main goal in addressing the geophysics environment is to enhance the awareness of the GRID community and the available services, in order to nourish a growing demand for utilizing GRID resources. Present achievements have been a promising contact to the geodetic environment, where a GRID – enabled method for Choleski factorization for equation solving, is under preparation.

France: The EGITF11 co-organized the EGITF11 with the EGI.eu team and the first France Grilles scientific grid day co-located are the main events of the year. Ten presentations, ten posters and several demonstrations based on work done by the French research teams were presented during the France Grilles scientific grid day. The France Grilles web site has been completely renewed in order to be more attractive for new user communities. Development of English pages are ongoing and will be finished in PY3. The France Grilles publications collection set up during 2011 is continuously growing and sums up to 182 articles and 278 references of publications now (74 for 2011 and 4 for 2012). This work will continue and statistics will be published soon. The CEFE laboratory that is the largest French research centre in Biodiversity is currently porting its first application on the grid. The BEDOFIH project that was selected within the framework of the "Excellence facilities" program launched by the French government aims to create a European intraday financial database. Its data will be stored in the France Grilles infrastructure (at the LPSC laboratory). CNRS, through Institut des Grilles et du Cloud and the LPSC laboratory is one of the BEDOFIH partners. France Grilles prepares now its Advisory committee meeting (April) and the France Grilles Day second edition. Several contacts with potential users are ongoing.

Georgia: Currently, has two groups are using grid infrastructure in Georgia:

- Meteorology: Investigation of Advanced Research WRF (ARW) modelling system for weather research and forecasting – Meteorology Department.

- Biophysical Chemistry: Modelling of some biochemical processes with the purpose of realization of their thin and purposeful synthesis – Tbilisi State University and Sokhumi State University.

NGI-GE supports users to adapt these applications from EGI AppDB to the national grid infrastructure and to make them aware about software upgrades. Regular assessment of research groups at universities and research institutes interested in participation GRI activities. GRENA together with research groups in Meteorology, Biophysical Chemistry, High Energy Physics and Plasma Physics are preparing project under the National Science Foundation for the development of grid infrastructure in Georgia.

Hungary: NGI_HU has launched the e-Science Café Roadshow event, which was held in November addressing topics such as Grids, Desktop Grids, Clouds, GPGPU and HPC. The Grid Application Support Centre team ported 4 new applications (KOPI, JChem, GATE, Conformer Generator), and one to Desktop Grid infrastructure is underway (Biome-BGC). Training activities included a DCI summer school 2011 and portal tutorials. New releases of WS-PGRADE/gUSE Portal (the latest is 3.4.1), with new features such as submitter monitoring, BOINC support in the submitter, uniform submission service (DCI Bridge) and statistics portlet on usage. Extending EGI with Desktop Grids middleware is

in good progress and is almost finished. The integration is conducted by the EDGI project team. The EDGI production infrastructure with more than 130k CPUs will be offered as a service for EGI and NGI user communities, thus enabling and exploiting the possibilities lying within volunteer and institutional desktop grid computing for a wider scientific user community. In PY3 the team aims to reach out to national ESFRI project leads and promote these existing services.

Ireland: Grid-Ireland as the Irish NGI has continued to work with existing grid users particularly in mathematics and astronomy, and incorporating research in making GPU resources available to users. Grid-Ireland continues to provide a support helpdesk for Irish grid users. Grid-Ireland has begun deployment of web portals (general purpose and application specific, e.g. for solar physics users) to replace a number of under-used command-line user interface services.

Italy: During the PY2 the NGI_IT user support activities have been reorganized and it now reflects the schema that is now in place in EGI-InSPIRE since PQ7. Outreach activities at a national level focused on organizing meetings with various communities in order to understand if their involvement in the national and European grid is possible. These communities cover numerous scientific institutes and projects, such as:

- The National Council of Research (CNR) about the D4Science and iMarine projects
- Italian Institute for Geophysics and Volcanology (INGV) about the ESFRI projects EMSO and EPOS
- Various bioinformatics groups connected to ESFRI projects ELIXIR, LIFEWATCH and BIOVEL
- The INFN experiments GERDA and ICARUS
- The Technological Transfer institute Mario Boella (New Generation Sequencing)
- Pharmacological institute Mario Negri (marionegri.it) (Matlab, OPENFOAM, R)
- LHCf experiment (the smallest of the six official LHC experiments)
- Astra Project (a the Science Gateway for the project was developed)
- INFN SPES experiment (porting and license handling of ANSYS packages)
- National Institute of Astrophysics (porting of VISIVO system)
- SPACI consortium and the ELETTRA synchrotron

The NGI_IT user support unit collaborated with other NGI_IT units in order to analyze the issues and improve the MPI and parallel jobs support of the Italian infrastructure, in particular the interest is in having a HTC/HPC integration as transparent as possible to the user. Dissemination activities were focused on reviewing and updating dissemination material and various events. The COMCHEM VO with NGI_IT support has setup an electronic magazine (ISSBN enabled) on which the VO Grid activities can be advertised.

Latvia: During PY2 Latvian NGI has actively participated in helping user communities to migrate workflows from gLite to ARC middleware. Initial effort was concentrated on creating middleware-agnostic workflows based on Condor job submission software, since Condor would be able to submit and manage jobs both in grid environment (ARC & UNICORE), as well as local cluster and user workstation levels. Due to several limitations of Condor, it is currently used only for some user front-

end needs, and due to lack of accounting interoperability with ARC is removed as a LRMS back-end from IMCS UL site. Currently, IMCS UL is evaluating use of SLURM LRMS. During PY3 both major sites in Latvian NGI plan to finally procure new hardware for increased computing resources. To harmonize middleware within Latvian NGI and NDGF_NGI, as well as allow easier use of national cluster resources, all sites in Latvia are migrating not only to ARC, but also to SLURM LRMS.

Lithuania: The team was very active in PY2. We took part in several scientific festivals where we reached out to new communities. We delivered courses on grid computing for students and provided consultancy to potential users. Leaflets, user guides about the national grid have been prepared. We would also appreciate more user guides solving common often repeating problem on grids and would have new materials for teaching and training. It will be preferable to have some "teaching standard". Difficulties that the NGI would like to solve in PY3 with support from EGI:

- Problems with a license policy of commercial software on EGI
- Very small user communities with different task and programs, so we have tray to find NGI that have similar user request to enlarge the specific user community.

Portugal and Spain (Ibergrid): The Portuguese NGI User Support Team reinforced the bond with the Spanish User Support Team through the establishment of a unique user support model serving both NGIs. The support model defines user support shifts strategies, common documentation and policies, and establishes the deployment of a central regional helpdesk serving users from both countries. During PY2, the Portuguese NGI User Support Team built submission frameworks for Portuguese researchers from Life Sciences (following research lines in comparative genomics and using Next Generation Sequencing techniques), and for a local community running an open source parser from Stanford University over English soap operas retrieved from the Gunterberg project. Simultaneously, it received a request from e-NMR VRC to enable access to the Portuguese infrastructure. Among the major PY2 highlights, we emphasize the IBERCLOUD initiative kick-off, aimed to provide a cloud infrastructure for IBERGRID users. In Spain the major contributions to the kick-off in the usage of the infrastructure in early 2012 in the general physics VO comes from two areas: astrophysics and simulations of Complex Systems, in particular spin glass simulations, both new users to the infrastructure. The Astrophysics users are studying the collisions between different isotopes of hydrogen in the atmosphere, and which are considered essential in the evolution of planetary atmospheres and interstellar areas. Currently they are submitting about 2000 computations per week, each of about 200 minutes duration. International VO users of Spain are major contributors to the areas of Computational Chemistry Fusion and Auger VOs. In this respect we want to remark the very positive experience of the Users Meeting that took place in Madrid in January 2012 the yearly IBERGRID conference, held in Santander, Spain; and the release of IBERGRID 2011 report.

Romania: The RO-01-ICI site has deployed Matlab with Distributed Computing Server toolkit for 16 workers in a private cluster and it is testing its integration with the grid middleware. RoGRID-NGI portal was improved and is used to disseminate information regarding events involving grid community and to support users. The operation team provides support for national resources like gridmosi.ici.ro VO and regional resources available in seegrid and see VOs. Besides the HEP

community, there are users of the EnviroGrids VO, and there is interest for the Globus middleware. For PY3 we plan to test Globus in addition of gLite middleware.

Serbia: Due to strong advertisement of AEGIS and EGI to the user communities in Serbia, a significant increase of AEGIS VO members (more than 20%) was recorded during PY2. The total number of members reached 112. Support to the Serbian chemistry community is continued in usage of NAMD, Firefly, and Gaussian applications. Software stack available at NGI_AEGIS grid sites is enriched with new chemistry, engineering, and mathematical software. AutoDock Vina application has been deployed on the request of our chemistry community, while the OpenFOAM (Open Field Operation and Manipulation) framework is ported by Serbian engineering community. In addition, in collaboration with the Mathematical Institute of the Serbian Academy of Sciences and Arts, ILOG CPLEX optimizer code has been ported to the AEGIS infrastructure. NGI_AEGIS Helpdesk and NGI_AEGIS website have been regularly updated. EGI Training Marketplace gadget was included to AEGIS website in order to help Serbian Grid users to find training material and events that meet their needs in a more efficient way. Beside these activities, central NGI_AEGIS user interface machine has been upgraded to Intel Xeon machine with 8 cores and 16 GB of RAM. In PY3 NGI_AEGIS will focus on organization of training events and improving of level of support to scientific communities. Plans for porting and deployment of new scientific applications to Serbian Grid infrastructure were already made with the Serbian researchers.

Slovakia: The team continued to support users concerned with the simulations of the spread of fire in tunnels, where the FDS (Fire Dynamics Simulator) model is applied (combined MPI with OpenMP and Parametric) using EMI-1 middleware functionality. Also users of WRF (Weather Research and Forecasting) system, chemical DIRAC application (Atomic and Molecular Direct Iterative Relativistic All-electron Calculations) and application from the electronics domain were supported. We started to collect user requirements in domain "Fire simulation" and "Speech in the grid" through the setup of related Virtual Team projects. We organized the "7th International Workshop on Grid Computing for Complex Problems" in Bratislava. Tutorial on Cloud computing was organized in the scope of GCCP2011.

Turkey: The team has been providing direct support to the national user communities from different disciplines for data transfer, software and tools installation, job submission and documentation through the common communication channels (e-mail, wiki and blog). The new user community settlements were also supported by presenting EGI infrastructure and abilities and providing grid certificates. The current situation could be understood that the national user communities need to be strongly oriented to the ESFRI applications. After the expansion of user community related to ESFRI applications, the community will support the integration, testing or usage of the applications to EGI. NGI_TR will in PY3 support user communities who deal with ESFRI applications. The leaflet indicating ESFRI applications and the efforts for the EGI infrastructure integration will be prepared and distributed through the national users to increase the awareness level. Dedicated slot for relation between ESFRI applications and EGI will be provided within the content of "National Grid and High Performance Computing Conference" that will be organized in 12-14 April 2012 Ankara-Turkey.

UK: The UK has led a strong program of community engagement and outreach during PY2. We held two one day workshops with our stakeholders, the community projects such as the ESFRIs and our

members, the UK's HEIs. These workshops were to gather requirements for shaping and driving the UK's Technical Roadmap. We held a further two workshops during a review by our funders that continued this process, the outcome of which is anticipated soon. The UK's SeIUCCR program of community engagement through its champion networks started in April2011 and we ran a highly successful training school for early career scientists in August. The residential one week school was very popular (120 applicants for 30 places) and covered usage of distributed computing infrastructures from local resources through to grids, clouds and HPC; software sustainability; and it introduced students to the NGI and EGI. SeIUCCR held a one day networking event at the UK's All Hands Meeting where it brought together community outreach activity from SeIUCCR, NGS, GridPP, SSI, XSEDE and EGI to learn from each other and share best practice. It is funding champions to attend the SSI's Collaborations Workshop in April 2012 and will hold a breakout session of new community engagement. The UK runs an effective support activity, receiving and answering approximately 600 user queries annually through its helpdesks. Excluding site issues, authentication/certificate issues are still the most predominant support request, followed by general queries, applications, using the UK's NGS cloud, and Authorisation/VOMS. The UK has managed to improve its user's experience with managing access to resources and reduce the burden on its support team in a number of ways during the past year. The SARoNGS (Shibboleth) access route has been implemented at more NGI resource centres, enabling users access via their institutional logon and shielding certificate management from the user. It has developed its Certificate Wizard allowing a browser independent way of applying for and managing certificates. It has further improved its website instructions for grid authentication and authorisation. Together these measures have seen a dramatic fall in the number of user queries the helpdesks have received.

2.3 Project Issues

2.3.1 Operations

Participation of NGIs to Staged Rollout activities has been expanding, but distribution of load across the partners was very inhomogeneous and in particular it did not always reflect the amount of effort committed in the project. The distribution of Early Adopters across deployed products is now being revised, as with the release of EMI 2 different platforms are being supported (SL5, SL6, Debian) and staged rollout effort needs to be concentrated on products whose deployment is requested at higher priority. A survey was conducted in March to collect information on priorities.

Several NGIs and Resource Centres have been frequently experiencing low performance in PY2, among these: NGI_ARMGRID (Armenia), NGI_BG (Bulgaria), NGI_ME (Montenegro), NGI_MARGI (Macedonia), NGI_RO (Romania), IGALC and a few sites in Asia Pacific. Problems differ considerably, among these the unreliable infrastructure of the machine room hosting resources and services, unavailability of expertise in running grid middleware, lack of experience with procedures and best practices. A support plan for some of the NGIs belonging to the South-East European region was discussed and approved by the project. GRNET is now supporting Armenia, Montenegro and FYRoM to overcome some identified technical issues, and training activities are being organized in collaboration with the technology providers.

Continued maintenance and third level support of software in EGI is paramount. Both will be challenged by the end of the two main projects that currently ensure provisioning of deployed software (EMI and IGE). The risks that are being faced are the discontinuation of maintenance and support of a subset of products, lower quality of the support that is currently subject to SLAs, phasing out of the external repositories, and a change in the software distribution processes that will have to be reflected with changes in the EGI software provisioning processes. These issues are in the TCB agenda. Information about products which need to be supported in the long term at higher priority, was collected in a survey in March 2012[146].

2.3.2 User Support

2.3.3 No issues were raised for this activity. Software Provisioning

During PY2 SA2 encountered a number of issues, of which the vast majority are technical issues that were managed at the activity level. Details about these issues described in more detail in the regular quarterly reports provided by the project. As reported in PQ5, the members of TSA2.3 have been consuming more effort than expected undertaking the quality criteria verification activities. The partners participating in TSA2.3, CSIC and LIP, have worked towards keeping their effort consumption under control. As described in the quarterly reports for PQ5 to PQ7, the actions taken to mitigate this overspending are showing the desired effect as

2.3.4 External Relations

No issues were raised for this activity.

2.4 Project Management

During PY2 a significant restructuring of the project took place. On the recommendation of the EC Reviewer's WP2 (NA2) and WP3 (NA3) was merged and national user support functions were transferred to WP4 (SA1) with no increase in the effort allocated to the work package. To enable the resources in the merged NA2 to be used effectively with minimum disruption then were organised into four main areas: central teams at EGI.eu (covering marketing and communication, strategic planning and policy development, community outreach and technical outreach to new communities), dedicated teams running community services (services established under NA3 relating to the Training Marketplace, Applications Database, Client Relationship Management system), coordination of NGI non-operational activities through the NGI International Liaison (NIL) role, and technical staff within the NGI that could be flexibility deployed within identified Virtual Teams (VT). The VTs represent an approach that allows experts interested in an identified problem that needs to be solved.

To support the management of the project the Project Management Board met 7 times. This was more frequent than planned and reflected the increased time spent on the agenda discussing strategic issues in its role as the project' strategy board – role identified in the EC Review. The EGI-InSPIRE Collaboration Board met twice to inform all of the partners on the project's progress and to approve the changes to the DoW being proposed by the Project Management Board.

[146] https://wiki.egi.eu/wiki/EGI_Operations_Surveys

The deliverables and milestones produced by the project still suffered from occasional delays, however these delays were less and more consistent than in PY1. These have been identified as coming from documents entering late into the review phase of the document production – the review phase itself was taking about 4 weeks as planned.

The deliverables and milestones proceeded as planned. D6.4 (Services offered by the Heavy User Community to other user communities) produced in PY2 is an update to D6.1 produced in PY1. These services have not changed a great deal during the last year and, therefore, the services reported in D6.1 remain current. Consequently, the decision was taken that after a slight update to D6.1 to deliver D6.4, the resources previously allocated to the production of a new D6.4 would be diverted to the achievement of Milestones 610 and 611.

2.4.1 Project Management Metrics

The production of deliverables and milestones has been much more predictable in PY2 and PY1 but has on average been delayed about a month. Analysis of these delays shows that the review process is taking the expected time and it is the entry of the report into the review phase that needs to be improved. Document authors/editors will be made aware of this and instructed to start preparing rafts earlier – especially documents that rely on collecting input from the NGIs – as some of this input only appears after multiple requests. The project management will also consider more frequently pushing ahead with a document that is lacking complete community input (as it has not been given when requested) in order to keep to the planned schedule, rather than delaying a document so that it has input from every NGI.

2.4.2 Coordination Activities

Members of NA1 and NA2 have attended the e-Consultation meeting in Lyon during the EGI Technical Forum, as well as a number of workshops held by the EC seeking input from the community (including EGI) into the Horizon 2020 planning cycle. Additional coordination activities took place through some of the EC funded projects with which EGI-InSPIRE has collaborations – as described in the following section.

2.4.3 Cooperation with Other Projects

During PY2, the project consolidated the collaboration with the 7 partners engaged during PY1 through MoUs and expanded considerably the strategic relationships with other activities through 14 more MoUs in the areas of resource infrastructure provider, virtual research collaborations and supporting projects.

The four technology providers (EMI, IGE, SAGA and StratusLab) provided regular input on the Unified Middleware Distribution and standards roadmap. Furthermore, the EMI project contributed both the EMI-1 and update releases into EGI's software provisioning process; while the IGE project properly integrated into the quality verification process and staged rollout for Globus components although some issue in the area of documentation still arise; the SAGA project supported the inclusion of the SAGA plugin through EMI and IGE to offer a client API for user applications. The StratusLab project

regularly contributed to the organisation of EGI technical and community forums through dissemination, training or technical sessions.

The collaboration with the HMRC community has not developed meaningfully as defined in the plan. The WeNMR project provided information about the related VRC, successfully integrated with the EGI helpdesk, provided requirements for the applications and services needed by the supported VRC, regularly attended the UCB and contributed to the update and prioritization of requirements. Alexandre Bovin has also given a keynote presentation at the last EGICF 2012. WeNMR was also the first "stories from the grid" video presentation. The WLCG project contributed to the User and General EGI Sustainability workshop in Amsterdam in January 2012 and through many of the operational structures set up within the project. LSGC representatives have taken part in almost a dozen UCB meetings and have actively participated in EGI major events. They have provided information about the area of expertise for the related VRC and described the services that can be shared in EGI and contributed to the requirement identification and prioritisation.

The MAPPER project provided some key documentation in the area of policy, review on applications, users, software and e-Infrastructures, and vertical integration plan. Consultancy has been provided by EGI on the integration of the relevant middleware components to enable the MAPPER use cases, some milestone have been delayed, while proof of concept of integrating PRACE and EGI are expected during PY3. The ScalaLife project provided training material and applications to be made available on the EGI User support services and by providing requirements from the EGI Life science community for ScalaLife.

The SIENA project organised a sessions at both the EGITF 2011 and EGICF 2012 related to standards roadmapping, EGI-InSPIRE attended CloudScape IV and also provided support and consultancy in the definition of the standards roadmap for Grids and Clouds.

The e-nventory project provided the context to define a number of NGI status indicators to track the evolution of grid infrastructures across Europe in terms of resources contributed and funding levels. The e.nventory portal provides interesting visualization capabilities of those indicators and allows their correlation with data from other pan-European infrastructures. Data was gathered from the early start of grid projects in Europe to document the growth of the infrastructure in terms of number of resource centres integrated and the amount of compute and storage resources provided.

In collaboration with the EDGI project, desktop grid services were integrated into the EGI service repository (GOCDB) and desktop grid probes – developed by EDGI – have been integrated with the Service Availability Monitoring (SAM) while the accounting integration is still in progress.

Together with the CHAIN project, a workshop was held in Amsterdam in January 2012 to discuss CHAIN integration roadmap and interoperations issues of the regional initiatives in developing countries with EGI. The implications on tools were also discussed. CHAIN regularly participated to EGI community events, and is facilitating the consolidation of sustainable grid infrastructures capable of integrate with EGI services. The CHAIN collaboration facilitated the establishment of a MoU with the South African Grid Initiative, which was signed in September 2011 during the EGITF 2012. SAGrid is

currently under integration into EGI alongside resources provided by the Ukrainian National Grid Initiative.

The joint work with SHIWA progressed according to the plans agreed in the MoU. The strong presence of SHIWA at EGITF 2011 and EGICF 2012, and the joint workshop on E-Science Workflows that EGI.eu jointly organised with SHIWA in February 2012 complete all, but one milestone. The SHIWA workflow repository is now linked to the EGI website (go.egi.eu/workflows) but it is still not connected to the EGI monitoring system. The reason is that SHIWA did not have effort to complete this task. As an extra collaboration not covered by the MoU SHIWA plans to develop a web gadget for this that could facilitate its uptake even further. (This depends on whether SHIWA can allocate someone for this.) The SHIWA project has been recently extended until the end of the summer, so the last milestone (report) will be completed later than was planned originally

The joint work with SCI-BUS started at the end of March, and so far there were no technical milestones, however joint work towards the MoU milestones are progressing, with activities in a recently started Virtual Team (Science gateway primer) and with a planned joint workshop at the technical forum. The ability to complete the development and integration work needed for the availability monitoring of the SCI-BUS Generic Gateway technology will largely depend on the Monitoring framework development activities in EGI-InSPIRE JRA1-SA1. The release of the monitoring infrastructure that could support the monitoring of community services in EGI has been already delayed several times.

The DECIDE project supported the definition of services needed and offered by the connected community. The integration of the helpdesk and the VRC integration plans are late with regards to the planned schedule.

The e-ScienceTalk project supported the communication activites of EGI-InSPIRE in various ways: joint booth at the eChallenges 2011 (Florence, Italy), contribution to the SC11 booth, managing media communication (iSGTW, GridCast, Twitter, flickr, youtube) for the EGI Technical Forum, common publications such as e-ScienceBriefings, EGI case studies dissemination.

The gSLM project contributed tutorial activities in both EGI Technical Forum and Community Forum in the area of IT service management, assessing SLM maturity, ISO Management of quality, services and information security according to ISO 9000, ISO/IEC 20000 & ISO/IEC 27000 - Overview and practical insights. The project also provided support in the review of the EGI-InSPIRE deliverable D2.18 “Evolving the EGI business model” and consultancy in the revision of the SLA/OLA in the context of ITSM best practices.

3 Deliverables and Milestones Tables

Id	Activity No	Deliverable / Milestone title	Original Delivery date(*)[147]	Revised delivery date(*)	Status (**)
D1.5	WP1	Quality Plan and Project Metrics https://documents.egi.eu/document/436	13	15	PMB approved
D2.9	WP2	Dissemination Plan https://documents.egi.eu/document/507	13	15	PMB approved
D2.10	WP2	EGI-InSPIRE Presentation https://documents.egi.eu/document/506	13	15	PMB approved
D3.3	WP3	User Community Support Process https://documents.egi.eu/document/661	15	17	PMB approved
D5.4	WP5	UMD Roadmap https://documents.egi.eu/document/612	15	17	PMB approved
MS111	WP1	Quarterly Report Template https://documents.egi.eu/document/513	13	13	PMB approved
MS506	WP5	EGI Software Repository- Architecture and Plans https://documents.egi.eu/document/503	13	14	PMB approved
MS609	WP6	HUC Contact Points and the Support Model https://documents.egi.eu/document/419	13	14	PMB approved
MS705	WP7	CIC Operations Portal Work Plan https://documents.egi.eu/document/525	13	15	PMB approved
MS706	WP7	Operational Tools Accounting Work Plan https://documents.egi.eu/document/531	13	16	PMB approved
MS217	WP2	Dissemination Handbook https://documents.egi.eu/document/541	14	15	PMB approved
MS306	WP3	User Support Contacts https://documents.egi.eu/document/564	14	15	PMB approved
MS409	WP4	Deploying Software into the EGI Production Infrastructure https://documents.egi.eu/document/478	14	15	PMB approved

[147] *(*) Dates are expressed in project month (1 to 48).*

(**) Status = Not started – In preparation – Pending internal review – PMB approved

*(***) Nature = **R** = Report **P** = Prototype **D** = Demonstrator **O** = Other, Deliverable id: for Milestone attached to a deliverable*

Id	Activity No	Deliverable / Milestone title	Original Delivery date(*)	Revised delivery date(*)	Status (**)
MS410	WP4	EGI Helpdesk and the NGI Support Units https://documents.egi.eu/document/522	14	16	PMB approved
MS411	WP4	Operational Level Agreements (OLAs) within the EGI production infrastructure https://documents.egi.eu/document/524	14	16	PMB approved
MS507	WP5	Deployed Middleware Support Unit Operations Procedures https://documents.egi.eu/document/504	14	15	PMB approved
MS508	WP5	Software Provisioning Process https://documents.egi.eu/document/505	14	15	PMB approved
MS707	WP7	Roadmap for the maintenance and development of the deployed operational tools https://documents.egi.eu/document/523	14	16	PMB approved
MS112	WP1	Quarterly Report https://documents.egi.eu/document/723	15	17	PMB approved
MS218	WP2	EGI Newsletter https://documents.egi.eu/document/724	15	15	PMB approved
MS219	WP2	Establishing the external relations area of the website covering the policy bodies and collaborating https://documents.egi.eu/document/544	15	15	PMB approved
MS307	WP3	User Support Metrics https://documents.egi.eu/document/675	15	16	PMB approved
MS412	WP4	Operational Security Procedures https://documents.egi.eu/document/649	15	16	PMB approved
MS610	WP6	Services for High Energy Physics https://documents.egi.eu/document/540	15	16	PMB approved
MS611	WP6	Services for the Life Science Community https://documents.egi.eu/document/683	15	17	PMB approved
D6.4	WP6	Capabilities offered by the HUCs to other communities https://documents.egi.eu/document/472	16	18	PMB approved
MS220	WP2	Review of the website content https://documents.egi.eu/document/601	16	16	PMB approved
MS413	WP4	Deployment plan for the distribution of operational tools to the NGIs/EIRO https://documents.egi.eu/document/704	16	17	PMB approved

Id	Activity No	Deliverable / Milestone title	Original Delivery date(*)	Revised delivery date(*)	Status (**)
MS414	WP4	Integrating resources into the EGI Production Infrastructure https://documents.egi.eu/document/650	16	18	PMB approved
MS509	WP5	Service Level Agreement with a Software Provider https://documents.egi.eu/document/615	16	16	PMB approved
MS612	WP6	HUC Software Roadmap https://documents.egi.eu/document/684	16	18	PMB approved
D2.12	WP2	Standards Roadmap https://documents.egi.eu/document/721	17	18	PMB approved
D4.3	WP4	EGI Operations Architecture https://documents.egi.eu/document/763	17	19	PMB approved
MS221	WP2	EGI Technical Forum	18	18	PMB approved
MS222	WP2	EGI Newsletter https://documents.egi.eu/document/878	18	18	PMB approved
MS113	WP1	Quarterly Report https://documents.egi.eu/document/881	19	21	PMB approved
D4.4	WP4	Security Risk Assessment of the EGI infrastructure https://documents.egi.eu/document/863	19	22	PMB approved
MS613	WP6	Training and dissemination event for all shared services and the other tasks within the activity https://documents.egi.eu/document/685	20	22	PMB approved
D1.6	WP1	Gender Action Plan https://documents.egi.eu/document/982	22	24	PMB approved
MS114	WP1	Quarterly Report https://documents.egi.eu/document/999	22	23	PMB approved
MS115	WP1	EGI Global Task Review https://documents.egi.eu/document/961	22	24	PMB approved
MS116	WP1	NGI International Task Review https://documents.egi.eu/document/963	22	24	PMB approved
MS117	WP1	Work of the Asia Pacific Region https://documents.egi.eu/document/1067	22	24	PMB approved

Id	Activity No	Deliverable / Milestone title	Original Delivery date(*)	Revised delivery date(*)	Status (**)
MS224	WP2	Security Activity within EGI https://documents.egi.eu/document/965	22	23	PMB approved
MS510	WP5	EGI Platforms Roadmap https://documents.egi.eu/document/970	22	24	PMB approved
MS614	WP6	HUC Software Roadmap https://documents.egi.eu/document/746	22	24	PMB approved
MS708	WP7	Operational Tools Roadmap https://documents.egi.eu/document/962	22	24	PMB approved
D1.7	WP1	Annual Report on Quality Status https://documents.egi.eu/document/1019	23	25	PMB approved
D4.5	WP4	Annual Report on the EGI Production Infrastructure https://documents.egi.eu/document/1059	23	26	PMB approved
D5.6	WP5	Annual Report on the status of Software Provisioning Activity and the Work of DMSU https://documents.egi.eu/document/1015	23	24	PMB approved
D6.6	WP6	Annual Report on the HUC Tools and Services https://documents.egi.eu/document/742	23	24	PMB approved
D7.2	WP7	Annual Report on Operational Tool Maintenance and Development Activity https://documents.egi.eu/document/1063	23	26	PMB approved
D2.14	WP2	Annual Report on EGI and its External Relations Activity https://documents.egi.eu/document/1069	23	25	PMB approved
D1.8	WP1	Annual Project Report https://documents.egi.eu/document/1144	24	25	PMB approved
D1.9	WP1	Quality Plan and Project Metrics https://documents.egi.eu/document/1071	24	26	PMB approved
D2.15	WP2	Marketing and Communication Plan https://documents.egi.eu/document/1070	24	25	PMB approved

Id	Activity No	Deliverable / Milestone title	Original Delivery date(*)	Revised delivery date(*)	Status (**)
D2.18	WP2	Evolving the EGI Business Model https://documents.egi.eu/document/1040	24	25	PMB approved
D2.30	WP2	EGI Strategic Plan https://documents.egi.eu/document/1098	24	25	PMB approved
D2.31	WP2	EGI Technical Roadmap https://documents.egi.eu/document/1094	24	25	PMB approved
MS226	WP2	EGI Community Forum https://documents.egi.eu/document/1107	24	24	PMB approved
MS418	WP4	Operational Level Agreements (OLAs) within the EGI production infrastructure https://documents.egi.eu/document/1057	24	26	PMB approved
MS513	WP5	Service Level Agreement with a Software Provider https://documents.egi.eu/document/1110	24	26	PMB approved

4 Explanation of the Use of Resources

4.1 Summary

Effort consumption over the whole project during PY2 remained on track with 101% of the overall effort.

Type	Work Package	Worked Hours Funded	Worked PM Funded	Committed PM	Achieved PM %
MGT	WP1	12,055.0	83.7	84.2	99%
COORD	WP2	36,421.6	262.4	294.2	89%
COORD	WP3	16,969.1	124.1	97.0	128%
SUPPORT	WP4	169,993.5	1,251.7	1,195.4	105%
SUPPORT	WP5	15,750.8	115.6	125.8	92%
SUPPORT	WP6	32,615.6	238.0	241.7	98%
RTD	WP7	11,065.8	82.4	95.1	87%
	Total	294,871.5	2,158.1	2,133.2	101%

There were variations in the effort reported between different WPs. WP1, WP4, WP5 and WP6 were effectively on target for the year. WP3 represents the effort reported during the first 6 months of PY2 when the WP3 existed as an independent WP. After PM19, WP3's effort was merged with the effort in WP2 and for the full year of activity the committed effort reaches only 89% of the expected committed effort. The merger of WP2 and WP3 caused some disruption to the reporting processes that had been established at the start of the project. Staff who had previously reported into WP3 now had to report into WP4 if the partner had 'spare' user support effort that they should report against, while the staff who now had to report into the new WP2 (either NILs or as part of a virtual team) had to be identified and entered into the timesheet system. As a result the effort that was committed in WP2 was reduced as these new working practices were established. WP7 remains on track with a profiling of effort in PY2 to gather requirements for accounting which will then be implemented in PY3 & PY4.

4.1.1 NA1

EGI.eu has achieved 96% of the worked PMs compared to the committed PMs. This ranges from 82% for TNA1.1 Activity management to 101% for TNA1.2E consortium management. This corresponds with the high level of activity required for the management of a large consortium, including project meetings, gathering end of year costs and administration for a large number of partners.

4.1.2 NA2 (merged with NA3)

For TNA2.1 and TNA2U.1E, the overall achieved PM is 110%, reflecting the additional effort needed for the management of the changes to the NA2 workpackage during PY2. TNA2.2E and TNA2U.2E

achieved 115% of committed PMs, again reflecting the high activity level required to bring on board the NILs into the outreach and marketing materials and to re-launch the website.

In PQ5 and PQ6, the effort reported by the NGIs into the TNA2.2 dissemination task varied widely when comparing committed to worked PMs, due to the small levels of effort involved when spread over individual quarters. However, the outreach effort from PQ7 was consolidated into the activities of the NILs, reported in TNA2.1N and TNA2.6N. Effort figures in these two tasks range from 0% for UPT, IIAP NAS RA, SIGMA and ARNES, to very low effort figures from IICT-BAS, CSC, IUCC, LIP, VR-SNIC and IMCS-UL. However, about 50% of NGIs are reporting at 80% or more into these two tasks, including UCY, CESNET, KIT-G, CSIC, CNRS, MTA KFKI, TCD, RENAM, IPB, STFC and UCPH. On the whole, these higher figures correspond to the high levels of engagement in the Virtual Teams, with the exception of UCPH which is reporting effort but has not engaged in any Virtual Team activity.

The AppDB developer team (IASA, NGI Greece) used less effort that was planned for PY2 (8.4PM from 11.8PM, 71%). The under spending compensated to the overspending of PY1 (14.6PM from 11.8PM, 124%) and resulted a 97% of resource use over the first two years of the project (23PM from 23.6PM).

The Training Marketplace developer team (STFC, NGI UK) used slightly more effort that was planned for PY2 (3.9PM instead of 3.6PM, 107%). The overspending was not significant and resulted a 105% overspending for STFC in the TNA2.5E activity for the first two years of the project (9.5PM instead of 9.4PM).

The VO Services and then the CRM developer team (CSIC and LIP, NGI Ibergrid) used significantly more effort that was committed for PY2 (9.1PM instead of 6,1PM, 149%). The overspending was caused by the change of focus of the group (from VO Services to CRM), which involved the transfer of VO services and related knowledge to SA1 and the overhead in the learning and setup of a new service. CSIC and LIP used 138% of the committed TNA2.5E effort in the first half of the project (16.7PM instead of 12.1PM). This overspending is expected to be compensated in the second half, because the routine operation and regular update of the CRM system are expected to require less commitment than the current level of spending.

4.1.3 SA1

Overall, the PMs recorded by SA1 amount to 102% of the committed effort, so the effort consumed is in-line with what committed by the project. The effort consumed for the EGI Global Tasks in SA1, amount to 91% of the expected effort. The PY2 under-spending compensates the overspending observed in PY1, which was due to the management of the operations transition from EGEE to a new set of independent national operations centres.

Three global tasks were affected by under-spending:

- coordination of operations integrations (TSA1.4)
- coordination of documentation (TSA1.8)

- provisioning of catch-all core services for emerging NGIs and user communities (TSA1.8)

Both documentation and integrations coordination tasks were affected by personnel turnover and tasks were reassigned according to the expertise available in the project consortium. Operations integration coordination was handed by KTH to SRCE as of PM19, while it is proposed to reassign the coordination of documentation from CSC to EGI.eu as of the start of PY3. Under-spending for the provisioning of core services was due to local personnel hiring policies that affected the partners delivering the task, however this problem did not affect the programme of work defined for PY2 and the services delivered.

The effort achieved for the delivery of NGI local tasks was not affected by any major deviation.

4.1.4 SA2

Effort consumption for TSA2.1 and TSA2.2 are within acceptable limits. The unspent effort in TSA2.2 may be used as a buffer for Quality Criteria that need to be developed for software components that may be added to the pipeline in PY3. Alternatively it may be transferred to TSA2.3 (within the same partner, 12A – CSIC) to partially mitigate the overspending in the Quality Criteria verification activities.

This is reflected in the effort figures for the partners in TSA2.3. The initial effort in PQ5 is very high, particularly for CSIC (1.5 times of the committed effort), which then continues to decrease over the course of PY2 down to equalling the committed effort in PQ8. This trend in reducing verification effort in PY2 may well be reversed in PY3 with the additional verification effort needed to include more components on more platforms from EMI-2 in a future UMD-2 release. It is expected that the extra effort needed to handle these additional components will be much less than for PQ5, and that the work in PY3 will see a much better alignment of committed and worked effort.

During PY1, for TSA2.4 GRNET showed very diverse effort spending among its JRU partners even though the balanced effort across the JRU was within limits of 24PM over the year. This was explained with different assignments in different phases in the project for different partners in the JRU. The effort reported for PY2 does not differentiate between individual members of the same JRU. Aggregated over all GRNET partners, the effort spent during PY2 consistently did not match the committed effort. Beginning with a slight under reporting of effort in PQ5, this discrepancy continues to grow to an effort gap of 2.2 PM in PQ8, totalling to 4.4 PM over PY2 for GRNET – 82% of committed effort. This is due to local budget constraints that restrict the hiring of staff to provide the matching local effort.

The DMSU (TSA.5) shows under reporting of effort for CESNET and INFN (70% and 82% effort consumption, respectively). While CESNET and INFN provide the most 2nd level support within the DMSU, KIT-G and VR-SNIC report nearly complete effort consumption for in fact very low numbers of GGUS tickets that fell into their support duties for UNICORE (KIT-G) and ARC (VR-SNIC). To mitigate this situation and to improve the effectiveness of the available resource, it is proposed that the DMSU will be merged with the TPM duties in TSA1.7 from PY3 for the remainder of the project.

4.1.5 SA3

Overall the activity delivered 97% of the expected effort across the work package. Within that there were a number of significant extremes:

- TSA3.1 - CERN is 201%: Due to additional activity management effort.
- TSA3.2.4 - CYFRONET is >300%
- TSA3.2.4 - CSIC is 149% which is balanced by an underspend by CISC in TSA3.2.5 by 85%.
- TSA3.3 – INFN is 0%
- TSA3.4 - EMBL still 0: Waiting for senior management changes within EBI/EMBL to take place to provide a focus for this activity which has been reflected in an updated DoW.
- TSA3.5 - INFN 35%
- TSA3.6 - CNRS 246%: Extra work due to the end of the HealthGrid Association which will be reflected by rebalancing of effort within the JRU.

4.1.6 JRA1

Overall the effort consumed by JRA1 in PY2 is 83%; 105% on WP7-E tasks (TJRA1.1 and TJRA1.2) and 68% on WP7-G tasks (TJRA1.3, TJRA1.4 and TJRA1.5). The slight under spending on WP7-G tasks is mostly due to the TJRA1.4 (accounting for new resource types) task which was spent at the level of 50%. This was expected since this task started in PY2 and the first phase of the work was devoted to requirement gathering and preliminary planning. Effort will be consumed in the next two year when the real development phase will start.

The TJRA1.3 under spending experienced in TJRA1.3 during PY1 has been reabsorbed in PY2 due to the extension of the task. Final TJRA1.3 effort consumption is 95% over the two years based on the initial PY1 allocation of effort.

The overspending reported by CSIC in TJRA1.2 (146%)is due to the deep restructuring of the tools maintained by the partner that implied a lot of development effort.

Other reported effort deviations are within acceptable limits and can be easily compensated in the coming years of the project.

4.2 Resources consumed per beneficiary

5 Financial Statements per Beneficiary

5.1 Summary

<<Provided by the PO>>

6 CERTIFICATES

<<Provided by the PO >>

7 Annex 1: Dissemination and Use

7.1 Main Project and Activity Meetings

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
11/5/ 2011	Amsterdam	UCB meeting	2	https://www.egi.eu/indico/conferenceDisplay.py?confId=473
9-10/6/ 2011	CERN	Federated identity system for scientific collaborations	8	Federated security infrastructures for various ES data infrastructure (Earth System Grid for Climate data, GENESI for satellite data…) not interoperable with the one of EGI and even among them. Discussions have been carried out for ES and several other domains. https://indico.cern.ch/event/129364
14//6/ 2011	Phone	UCB meeting	2	https://www.egi.eu/indico/conferenceDisplay.py?confId=500
30/6-1/7/ 2011	Amsterdam	First Review of EGI-InSPIRE	20	Formal official review of EGI-InSPIRE at the end of the 1st year of the project.
8/7/ 2011	University de Savoie – EDYTEM, Bourget du Lac	User meeting of MUST: computing and storage facility of Savoie University	-	http://indico.in2p3.fr/conferenceDisplay.py?confid=5353
15/7/ 2011	Phone	UCB meeting	1	https://www.egi.eu/indico/conferenceDisplay.py?confId=547
19-23/9/2011	Lyon	EFI-TF2011		https://www.egi.eu/indico/conferenceDisplay.py?ovw=True&confId=452
3/10/11-5/10/11	IPGP Paris, France	VERCE Kickoff	-	Collaboration with VERCE FP7 Seismology project
5/10/11	IPSL Paris, France	EGI/ESG meeting	-	Status of the actions, further actions, organization, time schedule.
24-26/01/2012	Amsterdam	EGI community sustainability workshop	2	https://www.egi.eu/indico/conferenceTimeTable.py?confId=709
23/11/2011	Amsterdam	TCB F2F	15	https://www.egi.eu/indico/conferenceDisplay.py?confId=672
27/01/2012	Amsterdam	CHAIN-EGI meeting	8	agenda.ct.infn.it/conferenceDisplay.py?confId=701

7.2 Conferences/Workshops Organised

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
5/5/ 2011	JINR, Laboratory of Information Technologies, Dubna, Russia	JINR Seminar on distributed computing and Grid-technologies	40	P.S. Berezovskiy (Keldysh Institute of Applied Mathematics RAS): "Job management in a grid with stand-alone resources" (Materials of Master's Thesis) http://lit.jinr.ru/view.php?lang=lat&var1=activ&var2=seminars&file=seminar_arch&menu=
9-14/5/ 2011	Beijing, China	Asia 3 2011 - Epikh School for Site Administrators	12	The aim of this tutorial was to prepare system administrators on the installation of grid sites and actually put together those facilities. The first days was dedicated to both the Condor and the gLite middleware services installation and configuration. The second part was dedicated to have the participants install sites, remotely (from the tutorial), at their home institutions (http://agenda.ct.infn.it/conferenceOtherViews.py?view=standard&confId=268).
12-13/5/ 2011	Amsterdam, The Netherlands	User Virtualisation Workshop	72	http://go.egi.eu/uvw1
13/5/ 2011	Jyväskylä, Finland	CSC 40 Years Seminar	15	Dissemination event about European resources available to Finnish researchers http://www.csc.fi/csc/kurssit/arkisto/csc40-jy
16-19/5/ 2011	Prague	TERENA 2011	525	https://tnc2011.terena.org/web/media/archive
16-20/5/ 2011	La Biodola (Pisa) Italy	CCR-Grid Workshop	120	http://agenda.infn.it/conferenceDisplay.py?confId=3530
16-26/5/ 2011	Beijing, China	Asia 3 2011 - Joint CHAIN/Epikh School for Application Porting	10	The Grid School aimed at increasing the number of scientific/industrial applications running on the EUChinaGRID infrastructure. The School also aimed at bringing new research groups to use the EUChinaGRID resources as well as to diversity the application domains and scientific communities (http://agenda.ct.infn.it/conferenceDisplay.py?confId=475).
23-25/5/ 2011	Ljubljana	4th Belle II Computing Workshop	40	http://kds.kek.jp/confLogin.py?returnURL=http%3A%2F%2Fkds.kek.jp%2FconferenceDisplay.py%3FconfId%3D6893&confId=6893

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
23-27/5/ 2011	Stellenbosch, Africa	Africa 5 2011 - Joint CHAIN/ SPECIAL/Epikh School for Grid Site Administrators	6	The aim of this tutorial was to prepare system administrators on the installation of grid sites and actually put together those facilities to increase the size of the SAGrid infrastructure. The first days was dedicated to the gLite middleware, services installation and configuration. The second part was dedicated to have the participants install sites, remotely (from the tutorial), at their home institutions (http://agenda.ct.infn.it/conferenceDisplay.py?confId=534).
30-1/5/ 2011	Helsinki, Finland	Computational methods in protein science workshop	20	Dissemination event about European resources available to Finnish researchers. Mr. Kimmo Mattila (CSC) gave a lecture concerning grid computing and guided a hands-on exercise about running virtual screening tasks in grid environment.
30/5/ 2011-9/6/ 2011	Stellenbosch, Africa	Africa 5 2011 - Joint CHAIN/ SPECIAL/Epikh School for Application Porting	16	The Grid School aimed at increasing the number of scientific/industrial applications running on the SAGrid production infrastructure. The School also aimed at bringing new research groups to use the SAGrid resources as well as to diversity the application domains and scientific communities (http://agenda.ct.infn.it/conferenceDisplay.py?confId=531).
31/5 -2/6/ 2011	JINR, Dubna, Russia	ATLAS Computing Technical Interchange Meeting	60	V.Ilyin (SINP MSU) Russian Grid Segment and Tier1 Perspectives V.Korenkov (JINR) Grid Activities in JINR A. Petrosyan, D.Oleynik (JINR) Dubna Tier 3 site. Current Status A. Petrosyan, D.Oleynik (JINR) Tier 3 Monitoring Project. Short term plans N.Kutouski(JINR) JINR infrastructure for Tier 3 simulation S. Belov, I. Kadochnikov (JINR) xRootd monitoring for local site. Tutorial.
2/6/2011	Singapore	1st Int. Workshop on Climate Change Data Challenges - (C2DC)	20	Conference URL: http://www.iccs-meeting.org/iccs2011/index.html Workshop URL: http://adm05.cmcc.it:8080/C2DC/Home.html

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
8-10/6/ 2011	Santander, Spain	5th Iberian Grid Infrastructure Conference	~ 70	- Co-coordination of the 5th Iberian Grid Infrastructure Conference and participation in the Scientific Advisory Committee and in the Editorial Board - Annual gathering discussion point for the DCI Iberian community for all activities from operations, applications and user support. Presentation of the major breakthroughs in the community - Conference URL: http://www.ibergrid.eu/2011/ - Programme URL: http://indico.ifca.es/indico/conferenceTimeTable.py?confId=358
9/6/2011	JINR, Laboratory of Information Technologies, Dubna, Russia	JINR Seminar on distributed computing and Grid-technologies	35	V.V. Galaktionov (JINR) "GridCom, GridCommander: graphical interface for Grid Jobs and Data management" http://www.jinr.ru/news_article.asp?n_id=934
20-23/6/ 2011	Santander, Spain	CMST workshop ICCSA 2011		The workshop (http://www.iccsa.org/sessions) was aimed at discussing user applications implemented on the EGI grid within the COMPCHEM activities. The discussion is extended to quality of users and quality of services
20-25/6/ 2011	Albena, Bulgaria	Special session:"HPC Grid Applications", during the 3rd AMITANS Conference	More than 90 participants.	1 tutorial and 9 contributed talks were presented demonstrating achievements of the Bulgarian Grid applications in the domains: environmental protections, computational fluid mechanics, semiconductors physics etc. http://2011.eac4amitans.org/2.html
24/6/ 2011	Oxford, UK	UK NGI Technical Roadmap Workshop	7 UK ESFRI project representatives and 4 NGI/EGI staff	
27-28/6/ 2011	Bristol, UK	HealthGrid Conference	50	http://bristol2011.healthgrid.org/
06/7/ 2011	Birmingham, UK	UK NGI Technical Roadmap Workshop	15 UK HEIs and 4 NGI/EGI staff	
11-13/7/ 2011	Hanburg, Germany	WLCG Collaboration Workshop	120	http://indico.desy.de/conferenceDisplay.py?confId=4019

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
11-16/7/ 2011	Budapest, Hungary	Joint European DCI Summer School 2011		http://indico.lpds.sztaki.hu/indico/conferenceDisplay.py?confId=10 http://www.lpds.sztaki.hu/eudciss2011/
19-20/8/ 2011	Frankfurt, Germany	Workshop "NMR applications on the Grid"	Around 20 participants	The workshop gives an introduction, overview and practical training on how to use some of the widely used and state of the art biomolecular NMR applications.
5-9/9/ 2011	Karlsruhe	Swiss Grid School 2011	80	http://www.swing-grid.ch/event/285525-swiss-grid-school-2011
5-9/9/ 2011	Karlsruhe, Germany	GridKa School	About 100 participants from different countries	Education in Grid, virtualization, storage, Cloud, user support; Grid and Cloud security, presentation of projects and achievements in Grid
7-9/9/ 2011	Savona, Italy	13th Plinius Conference on Mediterranean Storms, Disasters and Climate Change	-	Presentation of a set of 6 hydrological applications running on EGI and published in the Journal of Hydrology (June 2011) Interest of the FP7 EU DRIHM
11-24/9/ 2011	University Kinshasa, Democratic Republic of Congo	School on GIS, GPS and New Technologies (International Space Weather Initiative, University of Kinshasa)	-	Contribution: Organisation of a day about intensive computing (Grid, HPC...) Presentations of EGI & its applications, of first tests of Deskstop-Grid by RDC RDC project for Grid
12-15/9/2011	Abingdon, UK	Using e-Infrastructures for Research Summer School	27	http://www.ngs.ac.uk/seiuccr/summerschool Summer school for 30 students to learn about what tools are available to them and how to use them within the field of e-infrastructures. From local resources to grids, clouds and international e-infrastructures.
19/9/ 2011	Lyon, France	Premières rencontres scientifiques France Grilles	80	Presentation of the scientific work done in France with the help of the grid. First France Grilles prize awarded by the scientific committee http://france-grilles-2011.sciencesconf.org/
27/9/2011	York, UK	Meet the Champions workshop	~30	http://www.ngs.ac.uk/seiuccr/ahm2011workshop A workshop to promote e-infrastructure usage using "champions", researchers already strong in their usage to spread the word to other communities. Also a chance for different e-infrastructure projects to come together and learn from each other. Presentations from NGS, XSEDE and EGI. Held at UK All Hands Meeting, York.

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
12-13/10/ 2011	Poland, Poznan	e-IRG workshop	136	
13/10/ 2011	Prague	IT support at FZU	20	Overview of IT resources and support for the section I at Institute of Physics
13/10/ 2011	CC-CNES Toulouse, France	Grid and Cloud computing Seminar	-	Presentation of Earth Science applications and feedback
20-21/10/ 2011	Sofia, Bulgaria	2nd workshop Computing Centre "MADARA"	More than 70 participants	The workshop was organized by the Bulgarian third party in EGI-InSPIRE project – 5A-IOOCCP-BAS. 3 tutorial and 30 contributed talks were presented demonstrating achievements of the Bulgarian Grid applications in the Computational Chemistry domain. It was demonstrated how these applications run on BG08-MADARA and BG01-IPP Grid clusters. http://madara.orgchm.bas.bg/bg/
24/10/ 2011	Amsterdam, Netherlands	DCI Projects meeting	11	S. Andreozzi attended the meeting and took minutes https://www.egi.eu/indico/event/669
24-26/10/ 2011	II SAS, Bratislava, Slovakia	7th International Workshop on Grid Computing for Complex Problems (GCCP2011)	48	https://www.egi.eu/indico/conferenceDisplay.py?confId=652
25/10/ 2011	Prague	Campus Network Traffic Monitoring workshop	60	http://www.ces.net/doc/press/2011/pr111026.html
26/10/ 2011	KIT	GGUS Report Generator f2f meeting	20	https://www.egi.eu/indico/conferenceDisplay.py?confId=655
18 Nov 2011	Seattle	Gateway Computing Environments 2011 Workshop, in conjunction with SC 2011	-	http://sc11.supercomputing.org/ http://www.collab-ogce.org/gce11/index.php/Technical_Committee
14/11/2011	Seattle	WORKS'11 workshop on Workflows in support of Large-Scale science	75	http://works.cs.cardiff.ac.uk/index.php
10/11/2011	Amsterdam	NGI Int'l Liaisons Kick-off Meeting	35	https://www.egi.eu/indico/conferenceDisplay.py?ovw=True&confId=659
24-26/01/2011	Amsterdam	User and General EGI Sustainability Workshop	60	https://www.egi.eu/indico/conferenceDisplay.py?ovw=True&confId=709

Date	Location	Title	Participants	Outcome (Short report & Indico URL)
13th-14th December 2012	LIP, Lisbon, Portugal	1st Ibergrid Cloud Meeting	~15	Startup of the IberCloud initiative aiming to implement a cloud infrastructure in for Ibergrid Users. (http://ibergrid.lip.pt/OD/Others/1st%20Ibergrid%20Cloud%20Meeting)
Jan, 11, 2012	Bologna, Italy	Workshop on IGI portal - Data Management	6 NGI_IT staff + 14 people from IGI partner	Start-up meeting about data management through web portals, a one day workshop to overview the different technical solutions for the handling of the data storage for the grid/cloud user and applications portals. Several use cases have been presented by IGI partner and Italian communities. https://agenda.italiangrid.it/conferenceDisplay.py?confId=577
Dec, 13-14, 2011	Bologna, Italy	Security handling and monitoring in the INFN and IGI sites	27 (2 NGI_IT csirt + 25 site security officers)	This two days workshop has been organized together with the INFN security group, in order to review the following topics: 1) Incident response procedures adopted in EGI/IGI and national NREN 2) Tools and methods for the monitoring of site security 3) Tools and methods for the live incident response and post mortem analysis. https://agenda.italiangrid.it/conferenceOtherViews.py?view=standard&confId=480
16/1/ 2011	Ljubljana	24th EUGridPMA Meeting		http://agenda.nikhef.nl/conferenceTimeTable.py?confId=1644

7.3 Other Conferences/Workshops Attended

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
4-6/5/ 2011	Budapest	FET11		http://www.fet11.eu/
9-12/5/ 2011	Oslo	Nordugrid 2011	60	http://indico.hep.lu.se//conferenceTimeTable.py?confId=1047
11-13/5/ 2011	Prague, Czech Republic	22nd EUGridPMA Meeting	13	https://www.eugridpma.org/meetings/2011-05/
12/5/ 2011	Amsterdam	Cloud workshop	2	

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
16-19/5/ 2011	Prague	Terena Networking conference	500	https://tnc2011.terena.org/web/media/archive
14-16/5/ 2011	Zürich, Switze-rland	Documentation workshop	9 (2 from CSC)	9 (2 from CSC)
16-20/5/ 2011	Toulouse, France	CTA Meeting Toulouse		http://cta.irap.omp.eu/toulouse2011/
16-20/5/2011	Biodola, Elba island, Italy	Workshop CCR INFN GRID 2011		Coordination workshop with all INFN resource centers e and communities
17/5/ /2011	Lille, France	Lille Grid Day	45	Presentation of France Grilles (French NGI) and Life Sciences activities in France on grid infrastructure. Lille is a new site. http://indico.in2p3.fr/conferenceDisplay.py?ovw=True&confId=5324
17/5/ 2011	Sofia, Bulgaria	IPB at SEERA-EI meeting in Sofia	100	The Executive Agency "Electronic Communications Networks and Information Systems" of the Bulgarian government, as a part of SEERA –EI project activities, has organized a meeting on "HPC related policy and programs" for South Eastern Europe policy makers. This event is seen as a crucial step for planning the SEE common vision and strategy for eInfrastructures coordinated development and use. SCL's Antun Balaz present Serbian experiences, plans and expectations towards SEE high performance computing infrastructures development during the round table on "HPC initiatives in other SEE countries: Romania, Serbia, FYROM, Greece". More information and presentation is available at: http://www.scl.rs/news/743
18-20/5/2	Garching at	NGI-DE annual meeting	NGI-DE sites and others (D-	Talks from projects Initiative for Globus in Europe, D-Grid (including DGI-2) and EGI-InSPIRE; workshops on

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
011	Munich, Germany		Grid, IGE); about 80 participants from 12 countries	different topics such as support coordination, service monitoring, portals, security, accounting, GT4 migration, and others.
20-22/5/ 2011	Hisar, Bulgaria	Workshop for Supercomputer Applications	120	Liaison with national supercomputer applications project SuperCA++ , contacts with scientific advisors and young researchers, presentations and discussions about other HPC and Grid applications from domain of biomedicine and statistical physics.
20-24/5/ 2011	Santiago de Composte la, Spain	Pierre Auger Observatory Analysis Workshop	Jiri Chudoba	http://fpaxp1.usc.es/aanalysis/
23-24/5/2011	Berlin, Germany	ASPIRE Workshop	66	S. Newhouse provided a presentation and S. Andreozzi chaired a break out session to lead the discussion. A number of areas where identified to focus on new development of new Internet services for research and education
23-26/5/ 2011	Newport Beach, CA, USA	CCGRID	200	1. Presentation of the paper « The Grid Observatory » (full paper) (see below for refs) 2. Member of the PC 3. Member of the judging panel for 4th International Scalable Computing Challenge (SCALE 2011) Site http://www.ics.uci.edu/~ccgrid11/
26/5/ 2011	Lyon, France	GRISBI 2011 days	50	http://www.grisbio.fr/evenements/journee2011/programme/ presentation of GRISBI, work done and future work
26/5/ 2011	Cork, Ireland	Lero / EMC Seminar	~20, 1 from NGI_IE	Presented NGI_IE involvement in EGI and other DCI projects. Made contact with Lero, Irish Software Engineering Research Centre and EMC (http://emc.com) research staff with intent to collaborate in future.
26/5/ 2011	Sofia, Bulgaria	Jubilee Session " 50 years from the first computer centre in Bulgaria"	200	The President of Bulgaria and the Chairman of the Bulgarian Parliament gave speeches during the session. E. Atanassov presented the Bulgarian Grid Infrastructure and the contribution of the Bulgarian teams, researchers and developers of the Grid applications in the European e-infrastructure projects.
26-27/5/2011	Oslo	NOTUR2011	10	Norvegian HPC conference, where Norwegian EGI personell get to meet face to face.

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
30-31/5/ 2011	Strasbourg, France	LCG France	60	presentation of France Grilles (French NGI) http://indico.in2p3.fr/conferenceDisplay.py?confId=4941
31/5-1/6/ 2011	Paris, France	Green Days @ Paris	30	The goal of this 2-days workshop was to build up the French community interested in green computing. A working group on instrumentation should be created 4T 2011, and the community will reconvene in January 2012. Site http://perso.enslyon.fr/laurent.lefevre/greendaysparis/
7/6/ 2011	Santander, Spain	5ª Reunión Plenaria de la Red Española de e-Ciencia	150	http://www.e-ciencia.es/FichaEvento.jsp?externos=null&IDEvento=25
8-10/6/ 2011	London	IWSG-Life11	38	IWSG-Life'11 brought together scientists from the field of life sciences, bioinformatics and computer science. The aim is to exchange experience, formulate ideas and introduce up-to-date technological advances in molecular and systems biology in the context of Science Gateways (https://sites.google.com/a/staff.westminster.ac.uk/iwsg-life2011/scientific-programme).
8-10/6/ 2011	Santander, Spain	Ibergrid 2011	150	EGI-INSPIRE Software Quality Criteria: http://indico.ifca.es/indico/contributionDisplay.py?contribId=27&sessionId=11&confId=358 http://www.ibergrid.eu/2011/
9-10/6/ 2011	CERN	Federated identity system for scientific collaboration		Problem addressed: federated security infrastructures for various ES data infrastructure (Earth System Grid for Climate data, GENESI for satellite data...) not interoperable with the one of EGI and even among them. An ES Grid user needs to get a collection of certificates to carry out ES applications. Discussions have been carried out for ES and several other domains. https://indico.cern.ch/event/12936
10-6/ 2011	Bern, Switzerland	Swiss National Grid Association - Scientific Advisory Council 2011	20	http://www.swing-grid.ch/event/306650-swing-scientific-advisory-council-2011
14-16/6/ 2011	Zurich	Documentation Workshop	2 (total of 6 + 3 remote participants)	https://www.egi.eu/indico/conferenceTimeTable.py?confId=481#20110614

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
14/6/ 2011	Lille, France	Grid Day		http://www.lifl.fr/~touzet/calculintensif11.html
14/6/ 2011	Otranto (Italy)	CMCC Annual Meeting	100	http://www.cmcc.it/publications-meetings/meetings/otranto-13-15-june-2011-cmcc-annual-meeting
16-17/6/ 2011	Tbilisi, Georgia	The Regional Conference "THE WAY FORWARD FOR THE INFORMATION SOCIETY in the Eastern Europe and South Caucasus countries: Priorities and Challenges"	100	http://www.extend-ict.eu/closing-Conference
17-21/6/ 2011	Hamburg, Germany	ISC'2011	(10 from CSC)	Information about European resources available to researchers was distributed on the CSC stand.
24/6/ 2011	Paris, France	CNRS-CSIC meeting	10	
26/6-1/7/ 2011	JINR, University Centre, Dubna, Russia	Scientific school for Teachers from JINR member-states	41	S.Mitsyn (JINR) "What Grid is and how IT helps modern physics"
27/6/ 2011	Bristol, UK	HealthGrid	40	http://bristol2011.healthgrid.org/
27/6-8/7/ 2011	HTW Berlin	Grid educational event	50	Learning the Grid concept and job/data management on the Grids
31/6-1/7/ 2011	STFC RAL, UK	UK HEP SYSMAN meeting	~40, 7 from EGI	http://indico.cern.ch/conferenceDisplay.py?confId=145475

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
5/7/ 2011	JINR, Dubna, Russia	A session of the State Panel on high technology and innovations, devoted to development of scientific research infrastructures of facilities known as facilities of the "mega science" class, chaired by the Prime Minister of the Russian Federation V.V. Putin http://www.jinr.ru/news_article.asp?n_id=958	>100	Stand "Grid at JINR" presented and Google Earth dynamic overlays for WLCG and experiments shown See photo http://www.jinr.ru/img_news/11/080711/p1_b.jpg
4-8/7/ 2011	Stara Lesna, Slovakia	International Conference on mathematical Modeling and Computational Physics	78	V.V.Korenkov (JINR)"Grid Activity in Russia and JINR"
4-8/7/ 2011	Istanbul,Turkey	High Performance Computing and Simulation (HPCS-2011)		http://hpcs11.cisedu.info/program (run EGI booth during the conference as well as presented the accepted paper)
6/7/ 2011	Oxford	NGI Technical Roadmap Workshop for large projects	4 EGI people and 7 ESFRI/FP7 projects	
7/7/ 2011	Special Economic Zone "Dubna" Dubna, Russia	Student's scientific and technical school "Future personnel -2011"	>100	V.Mitsyn(JINR) excursion at the JINR Grid computing centre A.Uzhinsky(JINR) lecture "Grid activity in Russia and JINR"
8/7/ 2011	University de Savoie EDYTEM, Bourget du Lac, France	User meeting of MUST: computing and **storage facility of Savoie** **University**		http://indico.in2p3.fr/conferenceDisplay.py ?confId=5353
8/7/ 2011	Bologna, Italy	IGI Portal		Workshop on portals and scientific gateways

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
13-29/7/ 2011	JINR, Dubna, Russia	2011 Student Practice in JINR Fields of Research	70	Lecture V.V.Korenkov (JINR)"Laboratory of Information Technologies" Practice for students on grid-technologies and parallel computing
23/7/2011	Bologna, Italy	IGI Portal		Workshop on portals and scientific gateways
26/7/2011	Dublin, IE	Research Prioritisation TWG3 Technology and Social Media, Creative and Cultutral Enterprise		David O'Callaghan made a written submission on the importance of e-Infrastructure as an opportunity and setting out the NGIs strengths.
30.-31/8/ 2011	NDGF HQ, Kastrup, Denmark	NDGF All-Hands-meeting	31	https://portal.nordu.net/display/ndgfwiki/Meeting-2011-08-30+NDGF+All+Hands+Meeting
6-10/9/ 2011	Vienna, Austria	LMC8 -8th Liquid Matter Conference	2	Igor Stankovic and Milan Zezelj attended the LMC8 and presented their work during the poster sessions. The two posters were titled: 1. “Investigation of interplay between finite size scaling and aspect ratio in continuum percolating networks” by M. Zezelj, I. Stankovic, and A. Belic “Aggregation kinetics of short-range attractive particles: Brownian dynamics simulations vs. Smoluchowski equation” by I. Stankovic, A. Belic, and M. Zezelj http://lmc2011.univie.ac.at/
7-9/9/ 2011	Savona, Italy	13th Plinius Conference	M. Petitdidier	Presentation of a set of 6 hydrological applications-already running on EGI and published in the journal of Hydrology (June 2011) Outcome: Interest for the FP7 EU DRIHM
11-24/9/ 2011	Univ. Kinshasa, RDC	Space Weather School	M. Petitdidier	organisation of a day about intensive computing (Grid, HPC...), presentations of EGI & its applications, of first tests of Deskstop Grid by RDC - RDC project for Grid Outcome:Grid & HPC workshop planned
12-13/9/ 2011	Helsinki	IRISC 2011		http://irisc-workshop.org/irisc2011-helsinki/
12-14/9/ 2011	TOBB ETU, Ankara, Turkey	Molecular Dynamics Workshop II	Feyza Eryol	Researcher who runs molecular dynamics simulations were the participants of the national workshop. We had a presentation which introduced the NGI_TR, EGI and related activities.

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
15-16/9/2011	Belgrade, Serbia	TERENA TF-CPR workshop	40	Attended social media training session and networking meeting with NREN dissemination representatives. http://www.terena.org/activities/tf-cpr/pastmeetings/belgrade2011/agenda.html
20-22/9/2011	Sunny Beach, Bulgaria	Second Regional Conference " Supercomputing Scientific Applications and Industry Usage"	About 60 participants mainly from Bulgaria and SEE region	Liaison with the national representative in PRACE -1IP project, contacts with senior and young researchers, presentations and discussions about other HPC and Grid applications from domain of biomedicine, life Science and statistical physics. http://www.scc.acad.bg/
30/9/2011	Yerevan, Armenia	Workshop on National Grid Initiative in Armenia	~50	The status of the Armenian National Grid Initiative was presented and discussed.
1/10/2011	Belgrade, Serbia	TEDx Belgrade 2011 - Dare to Share	1	Professor Aleksandar Bogojevic gave a talk titled "Belgrade is the system"
5-9/10/2011	Kiev	ARC technical workshop	44	http://indico.hep.lu.se//conferenceDisplay.py?confId=1118
11-12/10/2011	Tennessee, USA	TAGPMA F2F meeting	15	Represented EGI and WLCG as relying party of the PKI (http://www.tagpma.org/node/59)
12-13/10/2011	Poznan, Poland	e-IRG workshop	136	S. Andreozzi attended the workshop and prepared a summary for the EGI blog. S. Newhouse presented and participated in the panel discussion. http://www.e-irg.eu/e-irg-events/e-irg-workshop-poznan-12-13-october.html
13/10/2011	CC-CNES Toulouse, France	Grid and cloud computing Seminar	M.Petitdidier (IPSL,FR)	Presentation Earth Science applications and feedbacks
14/10/2011	Phone conference	Discussion about Earth System Grid Federation and European Grid Initiative	IPSL:J. Raciazek & M. Petitdidier, P.Kershaw (STFC, UK),H. Schwichtenberg (SCAI, DE), G. Sipos (EGI)	Clarification of the ESGF security and discussion of possible solutions
23.29/10/2011	Barcelona	INFOCOMP 2011		http://www.iaria.org/conferences2011/INFOCOMP11.html
25/10/2011	Prague	Talks with IT professionals	100	http://www.cs.cas.cz/hsi/

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
26-28/10/2011	Florence, Italy	eChallenges 2011	150	Hosted a joint booth with e-ScienceTalk at the event. http://www.echallenges.org/e2011/default.asp?page=paper-repository
Nov 15, 2011	Livermore, CA (USA)	ASQ & GS-CAD Technical Talk	30	Technical presentations on Distributed and Grid Computing.
December 14, 2011	Dallas, Texas (USA)	PDCS2011 conference	20	Tutorial on GRelC http://www.iasted.org/conferences/speaker1-757.html
9/11/11	Munich, Germany	SLA4DGrid workshop	Michel Drescher	Keynote speaker for SLA in Cloud environments, https://documents.egi.eu/document/1003
6/12/11	Stockholm, Sweden	IEEE Workshop on Cloud interfaces and Virtualisation	Michel Drescher	https://documents.egi.eu/document/1004
2-3/11/2011	RAL, UK	2nd Workshop on Federated Identity system for Scientific Collaborations	60	D Kelsey was on organising committee, chaired one session and prepared and presented the Summary Vision Statement. Good progress on working towards common goal of federated identity. https://indico.cern.ch/conferenceDisplay.py?confId=157486
24/11/2011	Brussels	SciTech Europe	400	Steven Newhouse presented in a masterclass on EGI: linking computers across Europe for European Science http://www.publicserviceevents.co.uk/masterclass-client/187/scitech-europe/3639 (EGI was also a sponsor)
12/01/2012	Amsterdam	2012 Internet Society Event	250	Steven Newhouse served as a panel chair: http://isoc.nl/activ/2012-newyear.htm
30/01/2012	Brussels	ERA Conference	400	Post-event article in production. Event page at: http://ec.europa.eu/research/era/consultation/era-wrap-up-event_en.htm
16-18/12/2012	Ljubljana, SI	24th EUGridPMA	24	D. Groep (FOM): Review and update of the authentication profile guidelines; agreed new trust anchor release schedule for EGI; reviewed IGTF Wish List and EGI status, including a new adoption plan for secure hash algorithms in view of the status of deployed middleware. D. Kelsey (STFC): represented EGI and WLCG as a Relying Party. Led preparation of new Guidelines on Attribute Authority Service Provider Operations. V1.0 agreed during the meeting. Also made good progress on other normal IGTF business. For full details, please see my report at https://www.eugridpma.org/meetings/2012-01/summary.txt

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
13th-14th December 2012	LIP, Lisbon, Portugal	1st Ibergrid Cloud Meeting	~15	Startup of the IberCloud initiative aiming to implement a cloud infrastructure in for Ibergrid Users. (http://ibergrid.lip.pt/OD/Others/1st%20Ibergrid%20Cloud%20Meeting)
Jan, 11, 2012	Bologna, Italy	Workshop on IGI portal - Data Management	6 NGI_IT staff + 14 people from IGI partner	Start-up meeting about data management through web portals, a one day workshop to overview the different technical solutions for the handling of the data storage for the grid/cloud user and applications portals. Several use cases have been presented by IGI partner and Italian communities. https://agenda.italiangrid.it/conferenceDisplay.py?confId=577
Dec, 13-14, 2011	Bologna, Italy	Security handling and monitoring in the INFN and IGI sites	27 (2 NGI_IT csirt + 25 site security officers)	This two days workshop has been organized together with the INFN security group, in order to review the following topics: 1) Incident response procedures adopted in EGI/IGI and national NREN 2) Tools and methods for the monitoring of site security 3) Tools and methods for the live incident response and post mortem analysis. https://agenda.italiangrid.it/conferenceOtherViews.py?view=standard&confId=480
November 7-8, 2011	Bucharest, Romania	Eastern Europe Partnership Event - Policies for Development of E-Infrastructures in Eastern European Countries	100	To raise awareness of the importance of developing e-infrastructures in eastern Europe, TERENA organised the Eastern Europe Partnership Event on 7-8 November in Bucharest, Romania. Managers of national research and education network (NREN) organisations, research and development advisors and e-infrastructure experts discussed the enhancement of e-infrastructures and services in eastern Europe, and their further integration with pan-European e-infrastructure activities (http://www.terena.org/news/fullstory.php?news_id=3036) Presentations were made by representatives of multiple NGIs.

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
14.12. 2011.	University Computing Centre SRCE, Zagreb, Croatia	Third CRO NGI Day	50	Annual meeting of the Croatian grid community. The program this year was concentrated on aspects of cloud computing: Cloud Services for Croatian scientific and educational system; Scientific Cloud Computing infrastructure for Europe; Public collections of data. As in previous years, report on status and usage of the national grid infrastructure was given and grid community members presented their use cases, experiences and challenges with grid technology. http://www.cro-ngi.hr/dan/2011/?&L=0
13/12/ 2011	Tehran, Iran	4th HPC and Grid Workshop http://www.scl.rs/news/767	50	SCL's Vladimir Slavnic was an invited lecturer at the event. He lectured on e-Science and Grid computational infrastructures, introduced Grid projects (EGI-InSPIRE) and middleware topics. He held a series of hands-on sessions with colleagues using resources of Iranian Grid resource centers. Vladimir also worked together with Iranian colleagues on establishing new Grid services and making them available to researchers.
22-25/11/ 2011	Toulouse, France	Jres2011	~500	French NGI stand and the following contributions: presentations on IPV6 Care, server Virtualisaion, Quattorposter on medical imaging distributed computing
28 Nov - 29 Nov/ 2011	IFIN-HH, Magurele, Romania	RO-LCG 2011 Workshop	40	http://rolcg11.nipne.ro/presentations/JINRGRIDROM-Strizh-1.ppt http://rolcg11.nipne.ro/
16 Jan - 17 Jan/ 2011	JINR, Dubna, Russia	35th meeting of the Programme Advisory Committee for Condensed Matter Physics	50	Poster presentation http://indico.jinr.ru/conferenceDisplay.py?confId=280
23 Jan - 24 Jan/ 2011	JINR, Dubna, Russia	36th meeting of the Programme Advisory Committee for Particle Physics	50	Poster presentations http://indico.jinr.ru/conferenceDisplay.py?confId=279
30 Jan – 5 Feb/ 2011	JINR, Dubna, Russia	XIX International Conference Mathematics.Computer.Education. MKO-2012	350	Plenary lecture “Grid and Clouds in JINR and Russia” made by V.Korenkov, JINR

Date	Location	Title	Participants	Outcome (Short report & Document Server URL to presentations made)
2012-01-27	TCD, IE	What can grid and cloud computing do for you?	40	Seminar in School of Computer Science and Statistics, TCD. Outcome: Interest from TCD computer science researchers in using Grid-Ireland resources.
2012-01-30	TCD, IE	Grid and Cloud e-Infrastructures	10	Meeting with Intel Academic Research Programme representatives. Outcome: networking with a view to future collaboration
2.-3.11.	RAL. Oxford, UK	Federated Identity Workshop		
8-9/11/2011	Brussels, Belgium	European Gender Summit	200	Distributed an EGI postcard focusing on gender issues in the delegate pack, attended and blogged from the event. http://www.gender-summit.eu/index.php
24/11/2011	Brussels, Belgium	SciTech Europe	150	Hosted a booth in the exhibition area, delivered a masterclass on EGI, filmed for release on DVD http://www.publicserviceevents.co.uk/programme/187/scitech-europe
5-6/12/2011	Brussels, Belgium	Innovation Convention	500	Attended the event and blogged for GridCast blog http://ec.europa.eu/research/innovation-union/ic2011/index_en.cfm
5-8/12/2011	Stockholm, Sweden	eScience 2011	200	Attended the event, blogged for GridCast and published an article in iSGTW http://www.escience2011.org/

7.4 Publications

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Capabilities of the HPC Grid cluster at IICT-BAS	Automatica& Informatics		Accepted, 2011	E. T. A.	Atanassov Gurov Karaivanova
Framework for Service Composition in g-Lite	AIP Conference proceedings		Accepted, 2011		Radoslava Goranova
Atmospheric Composition of the Balkan Region and Bulgaria. Some Numerical xperiments	AIP Conference proceedings		Accepted, 2011	G. D. K. A. N. G. M.	Gadzhev, Syrakov, Ganev, Brandiyska, Miloshev, Georgiev, Prodanova
Quantum Critical Transport Near the Mott Transition	Phys. Rev. Lett.	10.1103 /PhysRevLett.107.026401	107 (2011) 026401	H. J. D. V.	Terletska, Vucicevic, Tanaskovic, Dobrosavljevic

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
The Role of Disorder on the Electronic Structure of Conjugated Polymers. the Case of Poly-2,5-bis(phenylethynyl)-1,3,4-thiadiazole	Phys. Chem. Chem. Phys.	10.1039/c1cp20329k		J. M. N. M.	Granadino-Roldan, Vukmirovic, Fernandez-Gomez Et al.
Nonlinear Bose-Einstein-condensate Dynamics Induced by a Harmonic Modulation of the S-wave Scattering Length	Phys. Rev. A	dx.doi.org/10.1103/PhysRevA.84.013618	84 (2011) Accepted	I. A. H. A.	Vidanovic, Balaz, Al-Jibbouri, Pelster
The DECIDE Science Gateway				V. A., R. B., T. C., G. L., S. M., F. P., R. R., D. S., I. C., A. S., D. P., L. L., K. J. B., F. T., J. G., B. V., G. F.	Ardizzone, Barbera, Calanducci, Fargetta, La Rocca, Monforte, Pistagna, Rotondo, Scardaci, Castiglioni, Schenoni, Perani, Leone, Blinowska, Turkheimer, Babiloni, Georges, Vellas and Frisoni
Grid technology For hydrological application	Journal of Hydrology,	DOI 10.1016 /j.jhydrol.2011.04.03	Vol. 403, 1-2, 186-199, 2011.	G. M. L. M. N. N. V.	Lecca Petitdidier Hluchy, Ivanovic, Kussul, Ray, Thieron
The Grid Observatory	In 11th IEEE/ACM Int. Symp. On Cluster, Cloud and Grid Computing, pp 114 – 123, May 2011		pp 114 – 123, May 2011	C. A. P. M. C. J. J. M.	Germain Renaud Cady, Gauron Jouvin, Loomis, Martyniak, Nauroy, Sebag

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Towards non-stationary Grid models.	Accepted for publication in The Journal of Grid Computing		9 :4, December 2011 (regular issue).	T. C. P. M.	Elteto , Germain Renaud, Bondon, Sebag.
The ClimateG testbed: towards large scale distributed data management for climate change.	Proceedings of the International Conference on computational science, 1-3 June 2011.			S. G. P M. H. S. J. A.	Fiore, Aloisio, Fox, Petitdidier,. Schwichtenberg, Denvil, Blower, Cofino
Recent improvements in HLRmon, an accounting portal suitable for national Grids	Proc. of the 2011 International Conference on Grid Computing and Applications	ISBN: 1-60132-181-3 pp. 10-13		M.	Savic
Optimization of intermolecular interaction potential energy parameters for Monte-Carlo and molecular dynamics simulations	Eighth International Conference on "Large Scale Scientific Computations", June 6-10, 2011, Sozopol, Bulgaria			D.	Shahpaski
Experience on running the Ibergrid infrastructure within EGI	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9			I. E. A.	Campos Fernandez Lopez
Fostering multi-scientific usage in the Iberian production infrastructure	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9			G. M. H.	Borges David Gomes
Towards Green Computing	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9			A. C. E.	Simon Fernandez Freire
EnergySaving Cluster experience in CETA-CIEMAT	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9			M.F. J.C. S.	Dolz Fernandez Iserte
Focusing on an integrated computing infrastructure: the IFCA experience	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. López García, P. Orviz Fernández, I. Cabrillo Bartolomé

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Orchestrating Services on a Public and Private Cloud Federation	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				J. Ejarque, J. Álvarez, H. Muñoz, et al.
COMPSs in the VENUS-C Platform: enabling e-Science applications on the Cloud	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				D. Lezzi, R. Rafanell, F. Lordan, et al.
Merging on-demand HPC resources from Amazon EC2 with the grid: a case study of a Xmipp application	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Lorca, J. Martín-Caro, R. Núñez-Ramírez, et al.
An Automated Cluster/Grid Task and Data Management System	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				L. Miranda, T. Sá, A. Pina, et al.
DISET protocol for DIRAC	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Casajús, R. Graciani
LHCb Grid resource usage in 2010 and beyond	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				R. Graciani, A. Casajus
Aggregated monitoring and automatic site exclusion of the ATLAS computing activities: the ATLAS Site Status Board	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				C. Borrego, A. Di Girolamo, X. Espinal, et al.
A Geographical Information System for wild ?re management	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Pina, A. Esteves, J. Puga, V. Oliveira]
Extending a desktop endoscopic capsule video analysis tool used by doctors with advanced computing resources	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				I. C. Oliveira, E. Dias, L. Alves, et al.
Exchanging Data for Breast Cancer Diagnosis on Heterogeneous Grid Platforms	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				D. Segrelles, J. M. Franco Valiente, R. Medina, et al.

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Analyzing Image Retrieval on Grids	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				O. D. Robles, P. Toharia, J. L. Bosque, et al.
The CHAIN Project	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				F. Ruggieri, A. Alberto, G. Andronico, et al.
Stellarator Optimization Using the Grid	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Gómez-Iglesias, F. Castejón, M. A. Vega-Rodríguez
OptiWeb: An optimization application for steel cut industries ported to the Grid in the framework of PireGrid project	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				J. Ibar, G. Ruiz, A. Tarancón, et al.
DataLight: data transfer and logging of large output applications in Grid environments	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				P. Abreu, R. Fonseca, L. O. Silva
Software Provision Process for EGI	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				M. David, G. Borges, J. Gomes, et al.
Support to MPI applications on the Grid	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				E. Fernández del Castillo]
EGI-InSPIRE Software Quality Criteria	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Simón, C. Fernández, I. Díaz, et al.
An Aspect-Oriented Approach to Fault-Tolerance in Grid Platforms	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				B. Medeiros, J. L. Sobral
A SLA-based Meta-Scheduling in Advance System to Provide QoS in Grid Environments	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				F. J. Conejero, L. Tomás, B. Caminero, et al.

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Vulnerability Assessment Enhancement for Middleware	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				J. Serrano, E. Heymann, E. Cesar, et al.
Web interface for generic grid jobs, Web4Grid	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Tugores, P. Colet
Population-Based Incremental Learning Algorithm to Search for Magic Squares	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				M. Cárdenas-Montes, J. M. Franco Valiente, A. Cortés Fácila, et al.
WRF4G: simplifying atmospheric modeling experiment design in distributed environments	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				V. Fernández-Quiruelas, L. Fita, J. Fernández, et al.
GSG Tool: General Scripts on Grid Creating generic scripts for problems on Astrophysics	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				J. R. Rodón, A. D. Benítez, M. Passas, et al.
Simulation of batch scheduling using real production-ready software tools	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Lucero
Analysis of Xen effciency in Grid environments for scientific computing	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				A. Tugores, P. Colet
Datacenters infrastructures remote management: a practical approach	Proceedings of the 5th IBERGRID. ISBN: 978-84-9745-884-9				E. de Andrés, A. Fuentes, T. de Miguel
From EGEE OPerations Portal towards EGI OPerations Portal	ISGC 2010, Taiwan. Proceedings	S.C. Lin et E. Yen "Data Driven e-Science: Use Cases and Successful Applications od Distributed Infrastructures", Springer pp. 129-140			H.Cordier, C.L'Orphelin, S.Reynaud, O.Lequeux, S.Loikkanen, P.Veyre

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
A federated system for sharing and reuse of images and image processing tools in neuroimaging	Computer Assisted Radiology and Surgery	CARS'11	Berlin, Germany, June 2011	B F C F	Gibaud, Ahmad, Barillot, Michel et al.
NeuroLOG: A framework for the sharing and reuse of distributed tools and data in neuroimaging	Organization for Human Brain	OHB'11	Québec city, Canada, June 2011	M M F C	Dojat Pélégrini-Issac Ahmad Barillot et al.
g-INFO portal: a solution to monitor Influenza A on the Grid for non-Grid users	HealthGrid'11	HG'11	Bristol, UK, June 2011	T-T Q-M T-H H-P	Doan Dao Vu Pham et al.
Multi-infrastructure workflow execution for medical simulation in the Virtual Imaging Platform	HealthGrid'11	HG'11	Bristol, UK, June 2011	R S B V	Ferreira Da Silva Camarasu-Pop Grenier Hamar et al.
Studies of Resistive Wall Mode stability in multi-parametric space using Grid infrastructure	38th European Physical Society (EPS) conference on Plasma Physics			D T M M P	Yadykin Żok Plóciennik Owsiak Strand
"The Climate-G Portal: The context, key features and a multi-dimensional analysis",	Future Generation Computer System	doi:10.1016/j.future.2011.05.015		S A G	Fiore Negro Aloisio
"Grid and Cloud Database Management"		Springer, ISBN: 978-3-642-20044-1, 1st Edition., 2011, 353 p. [available online on July 31, 2011]		S G	Fiore Aloisio (editors)
Influence of Interface Roughness Scattering on Output Characteristics of GaAs/AlGaAs Quantum Cascade Laser in a Magnetic Field	J. Phys. D: Appl. Phys.	10.1088/0022-3727/44/32/325105	44 (2011) 325105		M. Zezelj, V. Milanovic, J. Radovanovic and I. E. Stankovic

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Phase Diagram, Energy Scales, and Nonlocal Correlations in the Anderson Lattice Model	Phys. Rev. B	10.1103/PhysRevB.84.115105	84 (2011) 115105		D. Tanaskovic, K. Haule, G. Kotliar and V. Dobrosavljevic
Symmetry Reduction in Multiband Hamiltonians for Semiconductor Quantum Dots: the Role of Interfaces and Higher Energy Bands	J. Appl. Phys.	10.1063/1.3631048	110 (2011) 053710		S. Tomic, N. Vukmirovic
Charge Transport in a Quantum Dot Supercrystal	J. Phys. Chem. C	10.1021/jp206526s	115 (2011) 21409		I. H. Chu, M. Radulaski, N. Vukmirovic, H. P. Cheng and L. W. Wang
SPEEDUP Code for Calculation of Transition Amplitudes Via the Effective Action Approach	Commun. Comput. Phys.	10.4208/cicp.131210.180411a	11 (2012) 739		A. Balaz, I. Vidanovic, D. Stojiljkovic, D. Vudragovic, A. Belic and A. Bogojevic
Grid Interoperability Based on a Formal Design	Journal of Grid Computing	10.1007/s10723-011-9198-8	Volume 9 / 2011 , online first: http://www.springerlink.com/content/583056248h0w62r6/		Z. Farkas
Fostering the use of the Mediterranean e-Infrastructure with Science Gateways and Identity Federations	EUMEDGRID-Support User Forum (Lyon), France, 22-23 Sept. 2011			V. A., R. B., R. B., A. C., M. F., E. I., L. G., S. M., F. P., R. R., R. and D. S.,	Ardizzone, Barbera, Bruno, Calanducci, Fargetta, Ingrà, La Rocca, Monforte, Pistagna, Ricceri, Rotondo and Scardaci

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Parallel Computing in EGI	7th International Workshop on Grid Computing for Complex Problems (GCCP2011)			1. Viera 2. Miroslav 3. Peter	Sipkova Dobrucky Slizik
The Climate-G Portal: The context, key features and a multi-dimensional analysis	Future Generation Computer Systems	doi:10.1016/j.future.2011.05.015	Vol. 28: 1-8 (2012)	S. A. G.	Fiore Negro Aloisio
SPEEDUP Code for Calculation of Transition Amplitudes Via the Effective Action Approach	Commun. Comput. Phys.	11 (2012) 739 doi: 10.4208/cicp.131210.180411a arxiv: 1105.0542			1. A. Balaz 2. I. Vidanovic 3. D. Stojiljkovic Et al.
Charge Transport in a Quantum Dot Supercrystal	J. Phys. Chem.	C 115 (2011) 21409 doi: 10.1021/jp206526s			1. I. H. Chu 2. M. Radulaski 3. N. Vukmirovic Et al.
Symmetry Reduction in Multiband Hamiltonians for Semiconductor Quantum Dots: the Role of Interfaces and Higher Energy Bands	J. Appl. Phys.	110 (2011) 053710 Selected for Research Highlights in J. Appl. Phys. + 3rd most downloaded J. Appl. Phys. paper in September 2011 doi: 10.1063/1.3631048			1. S. Tomic 2. N. Vukmirovic
Phase Diagram, Energy Scales, and Nonlocal Correlations in the Anderson Lattice Model	Phys. Rev. B	84 (2011) 115105 doi: 10.1103/PhysRevB.84.115105 arxiv: 1106.4708			1. D. Tanaskovic 2. K. Haule 3. G. Kotliar 4. V. Dobrosavljevic
Relaxation Properties in a Diffusive Model of K-mers with Constrained Movements on a Triangular Lattice	Phys. Rev. E	84 (2011) 031109 doi: 10.1103/PhysRevE.84.031109			1. J. R. Scepanovic 2. I. Loncarevic 3. Lj. Budinski-Petkovic Et al.

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Influence of Interface Roughness Scattering on Output Characteristics of GaAs/AlGaAs Quantum Cascade Laser in a Magnetic Field	J. Phys. D: Appl. Phys.	44 (2011) 325105 doi: 10.1088/0022-3727/44/32/325105			1. M. Zezelj 2. V. Milanovic 3. J. Radovanovic 4. I. E. Stankovic
Globus Europe keynote	https://documents.egi.eu/document/901				
Cracow grid workshop keynote	https://documents.egi.eu/document/900				
http://www.iasted.org/conferences/speaker1-757.html	PDCS2011, 14 – 16, 2011, Dallas, USA				
Data Management	https://www.egi.eu/indico/sessionDisplay.py?sessionId=14&confId=452#20110919				
Portal Technologies	https://www.egi.eu/indico/sessionDisplay.py?sessionId=64&confId=452#20110920				
Cracow Grid Workshop keynote	https://documents.egi.eu/document/900				
Globus Europe keynote	https://documents.egi.eu/document/901				
The Climate-G Portal: The context, key features and a multi-dimensional analysis"	Future Geneeration Computer Systems	10.1016/j.future.2011.05.015.	Vol 28: P1-8 (2012)	S. A. G.	Fiore Negro Aloisio
SOMA2 - Open Source Framework for Molecular Modelling Workflows	Chemistry Central Journal		Vol 2 (Suppl 1): P4 (2008)	T. T. P.	Kinnunen Nyrönen Lehtovuori
SOMA - Workflow for Small Molecule Property Calculations on a Multiplatform Computing Grid	J. Chem. Inf. Model.		Vol 46 (Suppl 2): P620-625 (2006)	P. T.	Lehtovuori Nyrönen
Scientific Workflow Reuse through Conceptual Workflows	Workshop on Workflows in Support of Large-Scale Science (WORKS'11)		ACM, Seattle, WA, USA, 12--18 november 2011	N. J.	Cerezo Montagnat
IWIR: A Language Enabling Portability Across Grid Workflow Systems	Workshop on Workflows in Support of Large-Scale Science (WORKS'11)		ACM, Seattle, WA, USA, 12--18 november 2011	K. J. R.	Plankensteiner Montagnat Prodan

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Astronomical plates spectra extraction objectives and possible solutions implemented on digitized first byurakan survey (DFBS) images	Journal Information Theories and Applications		Vol. 18, Number 3, 2011, pp. 243-247, 2011		Aram Knyazyan, Areg Mickaelian, Hrachya Astsatryan,
Finite-size Scaling in Asymmetric Systems of Percolating Sticks	Phys. Rev. E	10.1103/PhysRevE.85.021101	85 (2012) 021101		M. Zezelj, I. Stankovic and A. Belic
E-Infrastructures for International Cooperation	Computational and Data Grids: Principles, Applications and Design	10.4018/978-1-61350-113-9.ch006	p. 141-193, IGI Global (2011)		G. Andronico, A. Balaz et al.
Numerical Simulations of Faraday Waves in Binary Bose-Einstein Condensates	Book of Abstracts, 12th International NTZ-Workshop on New Developments in Computational Physics		(2011) CompPhys11, p. 3		A. Balaz, A. I. Nicolin
The Green Computing Observatory :a data curation approach for green IT	International Conference on Cloud and Green Computing, Sydney : Australia (2011)		http://hal.inria.fr/inria-00632423/en/	Cécile Germain-Renaud, Frédéric Fürst, Michel Jouvin, Gilles Kassel, Julien Nauroy, Guillaume Philippon	
Characterizing E-Science File Access Behavior via Latent Dirichlet Allocation	4th IEEE International Conference on Utility and Cloud Computing (UCC 2011), Melbourne : Australia (2011)		http://hal.inria.fr/inria-00617914/en/	Kim Y., Germain-Renaud C.	
Towards Non-Stationary Grid Models	Journal of Grid Computing 9, 4 (2011) 423-440 -		http://hal.inria.fr/inria-00616279	Elteto T., Germain-Renaud C., Bondon P., Sebag M.	

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
Using Clouds to Scale Grid Resources: An Economic Model	Rapport de recherche (2011) 34 -		http://hal.inria.fr/hal-00657309	Rodero-Merino L., Caron E., Desprez F., Muresan A.	
Grid Activities at the Joint Institute for Nuclear Research	XXIII International Symposium on Nuclear Electronics and Computing Proceedings		68-73	S.D. Belov P. Dmitrienko V.V. Galaktionov et al	
Monitoring for GridNNN project	XXIII International Symposium on Nuclear Electronics and Computing Proceedings		74-76	S. Belov D. Oleynik A. Pertosyan	
VO Specific Data Browser for dCache	XXIII International Symposium on Nuclear Electronics and Computing Proceedings		145-147	M. Gavrilenko I. Gorbunov V. Korenkov Et al.	
RDMS CMS data processing and analysis workflow	XXIII International Symposium on Nuclear Electronics and Computing Proceedings		148-151	V. Gavrilov I. Golutvin V. Korenkov Et al.	
The Local Monitoring of ITEP GRID site	XXIII International Symposium on Nuclear Electronics and Computing Proceedings		186-191	Y. Lyublev M. Sokolov	
Distributed training and testing Grid infrastructure for JINR Member Sates	JINR Laboratory of Information Technologies, Scientific Report 2010-2011		22-23	Kutovskiy N.A	

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
JINR participation in the WLCG project during the 2010-2011 years	JINR Laboratory of Information Technologies, Scientific Report 2010-2011		24-26	Belov S.D. Galaktionov V.V. Gromova N.I. Et al.	
JINR Infrastructure for Tier3 Monitoring Development	JINR Laboratory of Information Technologies, Scientific Report 2010-2011		30-31	Kutovsky N.A. Kadochnikov I.S. Valova L. Et al.	
Tier 3 monitoring software suit (T3MON)	JINR Laboratory of Information Technologies, Scientific Report 2010-2011		31-34	Belov S.D. Kadochnikov I.S. Kutovsky N.A. Et al.	
Monitoring and accounting for GridNNN project	JINR Laboratory of Information Technologies, Scientific Report 2010-2011		68-73	Belov S.D. Korenkov V.V. Mitsyn S.V. Et al.	
Finite-size Scaling in Asymmetric Systems of Percolating Sticks	Phys. Rev. E	85 (2012) 021101; DOI: dx.doi.org/10.1103/PhysRevE.85.021101		1. M. Zezelj 2. I. Stankovic 3. A. Belic	

Publication title	Journal / Proceedings title	DOI code	Journal references	Authors *Initials*	Authors *Surname*
E-Infrastructures for International Cooperation	Computational and Data Grids: Principles, Applications and Design	p. 141-193, IGI Global (2011) Book chapter in Computational and Data Grids: Principles, Applications and Design, Ed. N. Preve, IGI Global (2011); DOI: dx.doi.org/10.4018/978-1-61350-113-9.ch006		1. G. Andronico 2. A. Balaz Et al.	
Numerical Simulations of Faraday Waves in Binary Bose-Einstein Condensates	Book of Abstracts, 12th International NTZ-Workshop on New Developments in Computational Physics	(2011) CompPhys11, p. 3; 24-26 November 2011, Leipzig, Germany		1. A. Balaz 2. A. I. Nicolin	
The Extension of TORQUE Scheduler Allowing the Use of Planning and Optimization Algorithms in Grids, v Cracow Grid Workshop,	Cracow Grid Workshop 2011			Klusáček Dalibor, Chlumský Václav, Ruda Miroslav	