EGI-InSPIRE

ANNUAL REPORT ON THE TOOLS AND SERVICES OF THE HEAVY USER COMMUNITIES

EU DELIVERABLE: D6.9

Document identifier:	EGI-doc-745- D6.9-FINAL
Date:	**14/05/2013**
Activity:	**SA3**
Lead Partner:	**CERN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/745

Abstract

This document is the 3rd and Final Annual Report on the HUC Tools and Services: a public report describing the current status of the tools and services provided within SA3 and their activity over the past year.



I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Jamie Shiers et al.	CERN/SA3	28/03/2013
Reviewed by:	Michel Drescher	EGI.eu	26/04/2013
	Afonso Duarte	ITQB-UNL	30/04/2013
	Stella Arnaouti	AUTH	30/04/2013
	Fotis Psomopoulos	CERTH-IAB	29/04/2013
	Ashiq Anjum	EGI.eu	01/05/2013
Moderated by:	Steven Newhouse	EGI.eu	07/05/2013
Approved by	**AMB/PMB**		14/05/2013

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	28/03/2013	Document for review	Mike Kenyon
2	30/04/2013	Response to reviews (those possible without iterating with SA3 partners).	Jamie Shiers

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/about/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 European Strategy Forum on Research Infrastructures (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:

- 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.
- The continued support of researchers within Europe and their international collaborators that are using the current production infrastructure.
- 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.
- Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects.
- Mechanisms to integrate existing infrastructure providers in Europe and around the world into a production infrastructure, so as to provide transparent access to all authorised users.
- 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

1 INTRODUCTION

This document is the Annual Report on the Heavy User Communities Tools and Services for Project Year 3 (PY3[1]) of the SA3 (WP6) work package of the EGI-InSPIRE Project. It is a public report that describes the current status of the tools and services provided within SA3 as well as the activity over the preceding year.

Heavy User Communities (HUCs) are Virtual Research Communities (VRCs) that have been using EGEE and now EGI routinely and thus have become more structured and advanced in terms of their Grid usage.

The communities identified as HUCs within this work package are:

- High Energy Physics (HEP)
- Life Sciences (LS)
- Astronomy and Astrophysics (A&A)
- Computational Chemistry and Materials Sciences and Technologies (CCMST)
- Earth Sciences (ES)
- Fusion (F)

Community	Description, capabilities and services offered
All HUCs (TSA3.2)	This task provides support for tools and services that are used by more than one community, including Dashboards, applications such as Ganga, Services such as HYDRA and GRelC, Workflows and Schedulers (SOMA2, Kepler, Taverna) and MPI.
High Energy Physics (TSA3.3) [MS611]	The High Energy Physics (HEP) HUC represents the 4 LHC experiments at CERN, which are fully relying on the use of grid computing for their offline data distribution, processing and analysis. Increasing focus is placed on common tools and solutions across these four large communities together with their re-use by other HEP experiments as well as numerous different disciplines and projects.
Life Sciences (TSA3.4) [MS611]	The Life Science (LS) HUC originates from the use of grid technology in the medical, biomedical and bioinformatics sectors in order to connect worldwide laboratories, share resources and ease the access to data in a secure and confidential way through health-grids.
Astronomy and Astrophysics (TSA3.5)	The A&A HUC is devoted to the evaluation of different solutions for the gridification of a rich variety of applications, as well as the accomplishment of a good level of interoperability among different technologies related to supercomputing, i.e. High Performance Computing and High Throughput Computing, Grid and Cloud.
Earth Sciences (TSA3.6)	Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting and many others. Their presence in SA3 is currently centred in the implementation, deployment and maintenance of the services needed to provide access from the grid to resources within the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR). The ES HUC includes also researchers and scientists working in the climate change

[1] Project Year 3 runs from May 2012 until April 2013 inclusive.

	domain. In particular most of them actively participate in the Climate-G use case. This use case exploits the GRelC service for distributed metadata management and the Climate-G portal as scientific gateway for this collaboration.

Table 1: Summary of the Heavy User Communities

2 SERVICES FOR HIGH ENERGY PHYSICS

2.1 Introduction

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.

Without doubt the high-point of PY3 was the announcement by ATLAS and CMS that, with great confidence, both experiments had observed a new particle in the mass region around 125-126 GeV; the Higgs Boson.

The fact that the worldwide grid computing infrastructure not only kept up with load and capacity demands that exceeded the plans but was also accompanied by a measurable improvement in service delivered is testimony to the efforts of the entire community. 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.

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.

The remainder of this section covers the tools and services that are further described in MS610 [MS610].

2.2 Analysis Tools and Support

2.2.1 HammerCloud

HammerCloud (HC)[2] 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.

HammerCloud has seen work on the backend that allows a more sustainable growth in service clients. The first action has been performed on the web service, improving the deployment with more machines for the HammerCloud cluster and deep optimizations in the code to make web views between 20% and, in some cases up to 3,000% faster. A new storage backend was also deployed to provide a high availability database cluster to improve data management and durability.

Additional test infrastructure was added and tighter integration achieved with development teams from the two LHC experiments, ATLAS and CMS. In the ATLAS case, integration with the Athena developers will allow nightly release testing on the Grid, improving deployment quality and reducing "hot fixes" after the deployment of new versions.

2.2.2 CMS CRAB Client

The CMS Remote Analysis Builder (CRAB) was the first analysis tool in CMS to aid users in configuring

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

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.

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 CRAB2 and to continue developing and commissioning CRAB3. CRAB2 developments included:

- Improvement of documentation to include the AsyncStageOut component. Amongst various changes, the two main development items were the automation of user data publication and a new monitoring system.
- A proposal that, during migration of the CrabServer to a new version of the REST APIs, support for data-publication on demand will be dropped and integrated into the AsyncStageOut functionality. AsyncStageOut currently takes care of data movement to the final storage element and is then responsible for the data injection into the dataset-bookkeeping system (DBS).
- Development of a new monitoring system, designed around the continuous replication functionality of CouchDB, was completed and tested by the integration team. The new system assumes that every distributed instance produces internal documents with summary information and that these documents are replicated in a central database at runtime; the AsyncStageOut tool implements a similar logic.

Development of CRAB3 has been completed with most of the basic functionality needed and version 3.1.4 has been released. Integration and beta testing were performed and useful feedback received. Key features added to CRAB3 during PY3 include:

- Support added to CRAB3 which allows users to generate small, private samples of Monte-Carlo data.
- Support of the input lumi-mask to enable the capability for the user to select the input data to be analysed at a finer granularity.
- Automate data publication through the AsyncStageOut service and the newly developed DBSPublisher component.
- The ability to perform a manual resubmission of failed jobs, respecting the security constraints.
- Other required functionalities to manage the workflow (to produce reports, monitor transfers and the publication status) and to perform troubleshooting in the event of failures (i.e. retrieve log file, kill pending jobs etc.)
- Improve web monitoring to track the progress of all workflows and in order to have an overview on the distributed system activities.

During PQ10 two distinct versions of the services providing these functionalities were released. In both cases there was intensive testing performed by the CMS Integration group, which included the participation of beta-users. In both test campaigns useful feedback was provided and a solution implemented in subsequent releases. Work has also included the refactoring of deployment scripts.

These were improved in order to automate the deployment of CRAB3 services on the CMS Cluster (cmsweb.cern.ch), allowing for the deployment of dedicated redundant services on which CRAB3 relies on.

Finally, CRAB3 components were evaluated with the PanDA server as a core system for the job life cycle management, replacing the WMAgent system. The study's main goal was to evaluate the future possibility of a common distributed analysis system for the ATLAS and CMS experiments. This study has led to a proof of concept prototype which has demonstrated to be able to run CMS jobs through the PanDA system already in production for the ATLAS experiment.

2.3 Data Management Tools and Support

2.3.1 ATLAS Distributed Data Management

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 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 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 and in any case it is outside the scope of EGI-InSPIRE SA3.

2.3.2 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 (more details have been given in document [MS610]). 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. The support of LHCbDIRAC, and in particular of its Data Management system (DMS), began in the EGI-InSPIRE project in October 2010.

The first preliminary version of the popularity service was put into production and evaluated. User-feedback and feature requests were implemented and deployed during PQ10. The popularity service provides metrics to quantify dataset popularity and provide a ranking of the most popular datasets (i.e. those most frequently accessed by users). The final goal is to use the information gathered to implement a dynamic data placement model, whereby the number of replicas of a given dataset is based on its popularity.

The LHCbDIRAC agent, which provides accounting plots for storage resources usage and which was first deployed in production during PY1, was refurbished and improved. The new implementation takes advantage of some new functionality in the framework, allowing more efficient usage of resources and a reduced number of queries to gather the necessary information. Following thorough validation, improvements were put into production during PQ10.

During PQ11, the popularity service was reviewed and then taken over by other developers, who committed to follow up with support and development of the service, according to the requirements

of the LHCb user community. Similarly, the service for storage resources accounting, extensively used for data management, has been documented and other data management experts will support this service.

The system for consistency checks between storage elements and file catalogue has been reviewed and documentation has been provided to the LHCb data management team.

The activity started during PQ10 aimed at optimizing the LHCb production management and reducing the man power needed to run the production system has been continued in collaboration with other members of the production team, who will take over and finish this work.

2.4 Monitoring and 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.

2.4.1 Experiment Dashboard

At the beginning of July 2012, CERN announced that the LHC experiments had observed a particle consistent with the long-sought Higgs boson. As has been mentioned several times during the CERN Higgs seminar, and in the following press conference, computing has been an integral part of this major achievement. The CERN Director General, Professor Rolf Dieter Heuer, presented a summary listing three essential tools: “the accelerator – the experiments – Grid computing".

One important component of the LHC Grid computing, the Experiment Dashboard, provides essential monitoring of LHC computing activities and has contributed to the success of the CERN scientific program. ATLAS presentations at the CERN Higgs seminar and ICHEP2012, a major particle physics conference held in Melbourne, Australia, included plots taken from the Experiment Dashboard illustrating the ATLAS job processing activity on Grid infrastructure (See Figure 1).

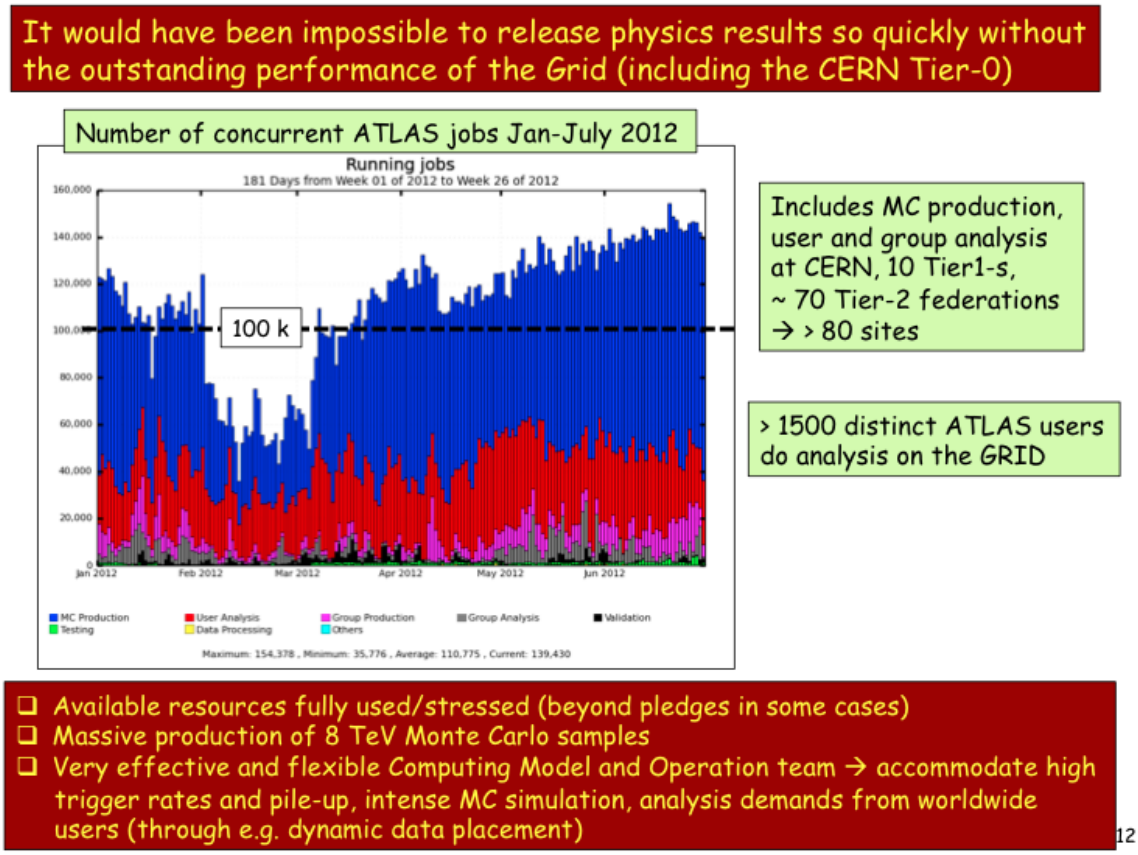

Figure 1. Slide from the ATLAS presentation at the CERN Higgs seminar displaying job-processing data from the Experiment Dashboard.

Experiment Dashboard applications have continued to be 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 the Dashboard SAM portal which provides information for evaluating site usability from the VO perspective.

A new version of the WLCG Transfer Dashboard was deployed in production. This version provides the ability to monitor not only data transfers handled by FTS, but also data transfers and data access performed by ATLAS and CMS on federated storage (xRootD). The next step is to deploy in production a version with the integrated xRootD traffic of the ALICE experiment. A WLCG Transfer Dashboard with ALICE xRootD traffic was prototyped and deployed to the integration server.

The Experiment Dashboard aims to provide a common solution for monitoring of the xrootd federations. Two prototypes with similar functionality but different persistency implementations are being developed. ORACLE is used as a database backend for the first prototype. Foreseeing a per-federation deployment model of the xrootd monitor, the Experiment Dashboard offers another solution with Hadoop/Hbase used for implementation of the monitoring data repository. The user interface, based on the xBrowser framework developed for transfer monitoring applications, is

shared by both prototypes and has a common core part with the WLCG Transfer Dashboard and ATLAS DDM Dashboard.

A new DDM accounting application to monitor the evolution of datasets, files and bytes over time was developed at the request of the ATLAS VO. The application provides advanced statistics either in historical or real-time views and it offers wide flexibility to the ATLAS users. The new application shares its implementation to a large extent with the Job Monitoring Historical Views, which allowed a prototype Proof of Concept (PoC) implementation to be developed in less than one month. The prototype is currently under validation by the ATLAS computing experts while feedback for additional features is constantly being received from the ATLAS experts. The final version of the application should be deployed in production in October 2012 during the ATLAS Software and Computing workshop.

A prototype Analysis Task monitoring tool, which includes the ability to kill jobs from the Task monitoring user interface, was deployed on the test server and is being intensively tested in order to make sure that user privileges are properly handled by the application.

The Site Usability Monitor (SUM) which provides visualization of the results of the remote tests submitted via the SAM/Nagios framework and site availability based on these results is heavily used by the LHC experiments for monitoring everyday operations. The data visualized in SUM is retrieved from the SAM repositories using the SAM APIs. Therefore validation of the new SAM releases should include validation of the SAM APIs. A set of tests to check the content and format of data retrieved with SAM APIs has been developed and is being used for validation of the new SAM releases.

Multiple improvements were performed in the Site Status Board application. Caching of data on the client side was implemented, which improves the performance of the user interface. The possibility to modify metric values from the user interface was enabled. The CMS production team evaluated SSB for resource usage monitoring. As a result of this evaluation a new production view was created in the CMS SSB instance. This view is now used by the CMS production operators. The SSB workshop was held for the SSB user community. During the workshop the SSB developers demonstrated new functionality of the application. SSB users from different LHC experiments shared their experience and provided feedback to the development team.

2.5 Summary of Services for HEP

As in previous years, the work of PY2 continued to focus on ensuring that the needs of the community in terms of production usage of the grid were met. As expected, the focus has changed from initial firefighting (PY1) to optimization (PY2) to production (PY3). Another key aspcct of PY3 has been to complete the task of handing over long-term support to teams that will continue long after EGI-InSPIRE SA3. This has been completed successfully, with the clear recognition from WLCG management boards that these activities are essential.

3 SERVICES FOR LIFE SCIENCE

3.1 Introduction

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. 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's production 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.

3.2 LSGC user management tools

The Life Sciences HUC steers the LSGC ("Life Sciences Grid Community" VRC[3]) 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 Communities 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[4];
- 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[5] for that purpose. New probes to monitor all VO SEs, WMSs and CEs were developed;
- On-line reporting tools easing the monitoring of SEs space management have been added to the technical team wiki page.

3.2.1 Work Accomplished

PQ9:

- New tools and web reports developed to allow for the monitoring of storage space consumed VRC-wise, and thus anticipate problems of storage resources running out of free space.

PQ10:

[3] LSGC wiki, http://lsgc.org/en/LSGC:home;jsessionid=C5A86D4A25C2EBA2EF62779A9D9B399D

[4] Biomed technical team wiki http://lsgc.org/en/Biomed-Shifts:Index

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

- Several Web gadgets customized for the Life Sciences have been added to the LSGC wiki, with the help of the User Communities Support Team [MS121].

PQ11:

- Upgrade of dedicated Nagios server to exploit new EMI-enabled probes and new topology builder based on VO feeds.
- Migration of the gLite VOMS server to EMI VOMS.
- Creation of an EMI UI virtual machine image made available for tests and deployment to the community.

3.2.2 Current work and Outlook

Besides the continuation of the current community support activity, discussions are continuing with Operations colleagues to improve and mutualise infrastructure monitoring tools and dashboards. The VO Administration and operation PORtal (VAPOR) is expected to help VOs mutualise their administration and operation effort, and save time for more community-specific activities. This development received financial support from EGI though its mini-project funding call, which will help finding the significant effort required in PY4.

3.3 GRelC data access interface

The GRelC service provides a WS-I compliant access interface to grid-databases. In this section, only the Life Sciences related activities performed during PY3 and regarding the GRelC interface are summarized, whereas a more comprehensive description of the results related to the GRelC software achieved during Y3 is presented in Section 7.1.

3.3.1 Work Accomplished in the LS area

During Y3, the following LS use cases have been supported: UNIPROT, Invasive Alien Species and GeneOntology. Some of them have been already finalized (such as UNIPROT, which started in Y1) whereas other ones (like the IAS) still need to be completed due to a different plan and schedule of the activities. During PY3, the support in terms of setup and maintenance of a GRelC node (deployed by the SPACI partner in Lecce) has been provided to address LS needs and host LS data banks like the ones related to the three use case mentioned before.

3.3.2 Current work and Outlook

The current (and future) work is mostly related to provide support in terms of new LS use cases definition and implementation as well as hosting of LS grid-databases. Moreover the porting activity of GRelC on EMI is an important milestone for this HUC too.

3.4 Hydra encryption service

Hydra 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.
Hydra has been officially released as part of EMI2. In addition, EMI has made significant efforts to produce useful documentation that was severely lacking. Nevertheless, the service remains hardly usable for production: some bug fixes are still on-going at EMI, while some concerns about the way

Hydra should be deployed and operated in a production environment have not been clarified yet. At present, the service that has been delivered remains a test service that gives the opportunity for the delivered functionality to be validated and the deployment procedures to be tested.
Finally, the decision of EMI not to make the Hydra client package mandatory in the Worker Nodes distribution suggests that VOs willing to use the service will have to either (i) deploy this package as VO-specific software on the computing elements they wish to use, or (ii) negotiate with each and every resource centre supporting the VO the deployment of the package. A survey has revealed that many production sites were misconfigured, not having deployed the Hydra client, having deployed an older version of the Hydra client, or publishing Hydra tags that are not consistent with the deployed client if any. A negotiation was led with each site publishing Hydra tags, or provisioning Hydra client without tags to resolve the situation.

3.4.1 Work Accomplished

PQ9:

- New Hydra key-store partner identified. Test service redeployed.

PQ10:

- Liaising with production-level sites either missing the Hydra client or hosting misconfigured or older versions.

PQ11:

- Hydra officially released as part of EMI2; significant efforts made to produce useful documentation.

3.4.2 Current work and Outlook

The production deployment model of the Hydra service remains unclear. Discussions are still ongoing with the EMI team developing Hydra, in order to figure out the best deployment model, and find answers to the concerns that remain unclear.

3.5 Issues and mitigation

Since the beginning of the project, the LSGC technical team has consumed most of its effort in performing basic monitoring of the grid resources and services accessible to the VRC, due to the complexity of middleware configuration and maintenance. Despite the automation of many monitoring procedures, the work remained dominated by the increasing number of incidents detected, which followed the increased capability of the team to identify misconfigured or faulty services. The team could hardly focus on application domain-specific tasks. Discussions are continuing with Operations to improve and mutualise infrastructure monitoring tools and dashboards.

Most sites only provide opportunistic access to their resources for this VRC. In this model, they may not allocate any real computing slots to the VOs in practice, although they will accept and queue biomedical computing tasks (jobs starvation). Initial discussions with EGI.eu, NGIs and largest site administrators to improve the VRC resources allocation policies have led EGI.eu to set up the Resource Allocation Task Force, which the LSGC technical team is a member of. Discussions on this topic are going on in the context of the EGI Community Forum 2013.

3.6 Summary

The Life Sciences HUC is a large, yet scattered community with fragmented user groups and very different use cases for grid infrastructure exploitation. In addition, the community is only supported to a limited extent by a large number of sites that only deliver best effort quality of services to the Life Science VOs. As a result, the HUC can hardly focus on the administration of one or a few community-specific services. It was engaged in a much more generic infrastructure monitoring and maintenance operation instead. This work is mostly domain-agnostic and could benefit to other user communities with similar profiles. The development of community administration tools continues until the end of the EGI project through the funding of the VAPOR mini-project.

4 SERVICES FOR ASTRONOMY AND ASTROPHYSICS

4.1 Overview

Activities carried out by the A&A community during PY3 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.
- Harvesting of astronomical workflows and applications to be ported on several distributed e-Infrastructures.

4.2 Coordination of the A&A Community

This activity, although not officially part of the work plan of task TSA3.5, is related to it given that one of the most important objectives of the task is the provision of requirements, use-cases and investigating: a) interactivity between e-Infrastructures based on different technologies (Grid, HCP and Cloud); b) access and management of astronomical databases from Grid Infrastructures. To fully achieve this, objective people working in task TSA3.5 (who are also in charge of the coordination of the European A&A community at large in EGI) need to interact with as many European A&A research groups and Institutes as possible.

The effort aimed at strengthening the presence of the community in EGI and to enhance the ability of the community to make use of DCIs continued during PY3. The majority of this activity was dedicated to the big transnational astronomical projects (especially the ESFRI projects) given their ability to attract new communities of end users. Although contacts with small projects and research groups are not neglected, priority was given to big projects in order to rapidly increment the number of end users and also for reasons related to the long-term sustainability. Contacts then were established with SKA, Euclid and CTA. Each of them acts as the reference projects for a specific branch of the astrophysical research (radio, astroparticle physics, etc.) with a strong ability to aggregate large fractions of the astro end users community. During PQ11 this activity mainly concerned the CTA ESFRI project with the creation of a VT (Virtual Team) for CTA and more in general for the whole Astro-Particle Physics community. The roadmap for the VT, finalized during the autumn of 2012, was submitted to EGI at the end of November 2012 and approved by EGI in early December. The VT has been running since January 7th 2013 and will be disbanded on July 7th 2013. The VT has been created to achieve the following objectives:

- Gather requirements from end users for what concerns Science Gateways and the SSO authentication system.
- Identify and put in place an identity federation model for the CTA collaboration and for the whole astroparticle physics community.
- According to the outcome of the user requirements gathering process, identify the most suitable technological solutions for the implementation of the SSO system and of one or more specialized Science Gateways.
- Define a roadmap for the design and implementation of the SSO system and of the Science Gateways after the completion of the VT activities. In the next (and last) quarter of EGI-

InSPIRE WP6 these activities then will be carried out and proposals for joint activities with other large astronomical projects (SKA, Euclid and others) will be formulated.

4.3 VisIVO

During PY3 significant results have been achieved for what concerns the porting to the Grid of VisIVO[6] (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. VisIVO-related activities competed during PY3 include:

- The study and porting of the VisIVO MPI version to the 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 movies could last several days. For this reason the MPI parallelized version of VisIVO plays a fundamental role. A parallel application for the Gaia Mission porting activity on grid gLite middleware has begun. The parallel application is dedicated to the development and test of the core part of the AVU-GSR (Astrometric Verification Unit - Global Sphere Reconstruction) software developed for the ESA GAIA Mission. The main goal of which is the production of a microarcsecond-level 5 parameters astrometric catalogue. The parallel code uses MPI and openMP (where available); it is characterized by an extremely low communication level between the processes, so that preliminary speed-up tests show behaviour close to the theoretical speed-up. Since AVU-GSR is very demanding on hardware resources, the typical execution environment is provided by Supercomputers, but the resources provided by IGI are very attractive for debugging purposes. Porting to EGI is in progress in the framework of the IGI HPC test-bed in which we select resources with a large amount of global memory and a high speed network, such as the one provided by INFN-PISA and UNI-NAPOLI sites. It is important that the use of the grid as an auxiliary resource provides us with the capability to make a large number of runs with a reduced number of stars as a validation step of the whole process.
- 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 therefore considered and a preliminary study focused on the porting and optimization of the data transfer between the CPU and GPUs on worker nodes where GPUs are available.
- The production of a CUDA-enabled version of VisIVO for gLite. A first preliminary study focused on the porting and optimization of the data transfer between the CPU and GPUs on worker nodes where GPUs are available. To provide a service able to take advantage of GPUs on the Grid, A&A acquired a new hybrid CPU-GPU system (funded by the Astrophysical Observatory of Catania) configured as a Grid computing node.
- The design and implementation of a specific grid-enabled library that allows users to interact with Grid computing and storage resources.

It is worth noting that the current version of VisIVO is also able to interface with and use the gLite Grid Catalogue and that, although VisIVO has been conceived and implemented as a visualization

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

tool for astronomy, recently it evolved in a generic multi-disciplinary service that can be used by any other community that needs 2D and 3D data visualization.

5 SERVICES FOR EARTH SCIENCE

5.1 Overview

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 Interoperations - 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.

5.2 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.

A new version of the GI-cat distributed catalogue service, which is able to broker between heterogeneous search and metadata infrastructures, has been deployed on SCAI's infrastructure. The Command Line Interface (CLI) and Text User Interface (TUI) gsearch, based on ncurses was adapted to handle this newer version, and provide a better overview of the different search results. Moreover, some misleading navigation errors were resolved and documentation improved. In order to limit the load on the different opensearch sites, gsearch now by default only requests a limited number of results. The user triggers the next request by hand or in the TUI case, by going to the last page of the results. In CLI mode, gsearch has a new batch system, which automatically downloads the requested resources and creates a job submit script. In addition, the user can provide a template containing the user application and settings, which will be used to generate the job submission script and is able to be directly submitted. There is ongoing effort to separate the opensearch specific parsing and request engine of gsearch into a dedicated development library. Third-party software or simulation applications could directly use the API to access data from opensearch sites and thus benefit from the development done by gsearch for a reliable access to ES data. Although many different cases and scenarios were tested, there are still sites which respond in an unexpected manner. These issues will be further addressed and resolved.

The web interfaces internal search was completely redesigned and now supports threading. Also, the full query at the catalogue sites is now made on user requests. So far, all GI-Cat attached catalogues have been immediately queried without user interaction. Now, based on preliminary information users can decide which site they want to query for full result information (including the file links etc.) The credential handling and job submission with jSAGA were extended and minor bugs have been fixed.

The VO for the VERCE project has been finalised (VERCE.eu) and the first two sites have joined. Deployment of VO software (ObsPy: seismological software) in the VO SW area of the sites has

started and first tests have been carried out. An official Memorandum of Understanding between EGI.eu and the VERCE EU project is under preparation. Members of the Earth Science Grid community have again arranged for a session at the General Assembly of the European Geosciences Union.

5.3 ESG

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 Earth System Grid (ESG) interoperability has made considerable progress. The intelligent data transfer tool, named “Synchro-data”, that facilitates the command line, bulk oriented access to ESGF data was updated. The tool can download files from the ESGF infrastructure 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 models. 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 potential new files.

Major added features include a new discovery engine, discrepancy detection and data version management.

The new discovery engine (Search-API) gives the ability to use projects other than CMIP5 (until now, Synchro-data was only able to download data from the CMIP5 archives). It is also faster than the previous discovery engine. The "Search-API" discovery engine works by calling a distributed server-side dedicated search API, while the former "THREDDS-catalogue" discovery engine works by parsing THREDDS XML catalogues , which is time consuming.

For the authentication interoperability, the prototype adaption of MyProxy has been released as CTS version 1. From the client side, CTS works as follows. First a "CTS patched MyProxy client" must be installed on an interactive machine (the patch add the "myproxy-bind" command to the myproxy distribution). When a user wants to access ESGF from EGI, he logons on that machine, creates an EGI proxy certificate with the voms-proxy-init command, then binds it with his ESGF account using the "myproxy-bind" command (ESGF password will be prompted at this step). From that moment on, the user can retrieve an ESGF short-lived certificate with the "myproxy-logon" command without entering a password, by just using the EGI proxy certificate. Thus, it is possible to access ESGF data from non-interactive EGI worker node (for better use, a script can wrap the "myproxy-logon" command to make the ESGF authentication step transparent).

On the server side the CTS patch must be installed on the ESGF MyProxy server. The patch makes the server able to process "myproxy-bind" requests, add a new mapping table for EGI<=>ESGF identities bindings, and modify "GET" requests handling the use the new mapping table.

Future plans for CTS include porting the patch to other MyProxy versions, at least 4.6 and 5.9 (HEAD), improving code quality, adding a new option to "myproxy-logon" command and improving the documentation.

6 OTHER COMMUNITIES & SHARED SERVICES AND TOOLS

6.1 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.

The most relevant activities related to this service that have been carried out during PY3 are:

1. new gLite3.2 compliant release of GRelC;
2. implementation plan for the GRelC service towards EMI1 and EMI2;
3. a new version of the DashboardDB registry and monitoring gadgets (permalinks available);
4. design and implementation of the GRelC Desktop;
5. dissemination activities in close collaboration with NA2.

Concerning point 1, a new release of GRelC for gLite 3.2 has been made available to the HUC during PY3. In this regard, a new version of the Globus libs (external libraries for GRelC) has been released for the SL5. The new rpm replaced the old one for the GRelC release on gLite 3.2 as well as it is needed by the EMI1/EMI2 compliant GRelC releases on SL5 x86_64. The GRelC rpms[7] and repository[8] (both available at IGI level) have been updated and tested accordingly.

Regarding point 2 (porting of GRelC on EMI1 and EMI2), a preliminary study regarding the compatibility of GRelC software with the EMI distribution has been carried out during Y3. After that, an implementation plan has been jointly defined with the IGI release group and some installation tests to port the GRelC software on EMI1 and EMI2 have been carried out. Consolidated results on the porting activity will be available by April 2013.

Concerning point 3, during PY3 the DashboardDB registry and monitoring gadgets have been refined, fixed and deployed twice. New community-based features and monitoring views have been added to the system. Four permalinks have been made created to export each gadget both in a secure (login-password based) and guest (free access) way into existing web applications just using a single HTML line of code. The permalink feature has been exploited to integrate the DashboardDB application into three different web-applications are importing the two gadgets (EGI Website[9], GRelC Website[10], GRelC Desktop application[11]). During Y3, the DashboardDB has proved to be a production-level application (0 errors detected in the last 6 months).

Regarding point 4, a web desktop application (called *GRelC Desktop*) including the two DashboardDB gadgets has been made available to the HUC during Y3. The GRelC Desktop is a flexible environment joining the pervasiveness and platform independence of a web-based application with a superior user experience and responsiveness related to a desktop-based application. It includes all of the gadgets implemented during the project and new ones related to well-known social networks like Twitter and Youtube.

Some examples are:

- the DashboardDB registry (both secured and guest-based);

[7] Rpm repository (3_2_0 - x86_64) http://repo-pd.italiangrid.it/mrepo/ig-cert_sl5-x86_64/RPMS.3_2_0/

[8] GRelC metapackage http://repo-pd.italiangrid.it/mrepo/ig-cert_sl5-x86_64/RPMS.3_2_0/ig_GRELC-3.2.0-0_ig26_sl5_test09.x86_64.rpm

[9] EGI Web Gadgets: http://www.egi.eu/services/support/gadgets/

[10] GRelC Website: http://www.grelc.unile.it

[11] GRelC Desktop Application: http://adm05.cmcc.it:8080/GrelcDesktop/

- the DashboardDB monitoring (from global to service based views);
- the Twitter gadget to follow the activities related to the DashboardDB application (the "DashboardDB" account has been created during Q10);
- the Youtube gadget for dissemination purposes.

The GRelC Desktop is very extensible, easy to use and new gadgets can be straightforwardly included as new "apps". Moreover the desktop approach allows keeping several "apps" active at the same time in separate windows (see Figure 2). It is important to remark that the GRelC Desktop provides both "secured" (through login/password) and "guest-based" gadgets (grid-certificates are not needed to carry out the authentication step). Finally, the GRelC Desktop aims at integrating in a web-desktop based environment all of the resources related to the GRelC software (GRelC website, DashboardDB gadgets, dissemination material, community-based gadgets, etc.).

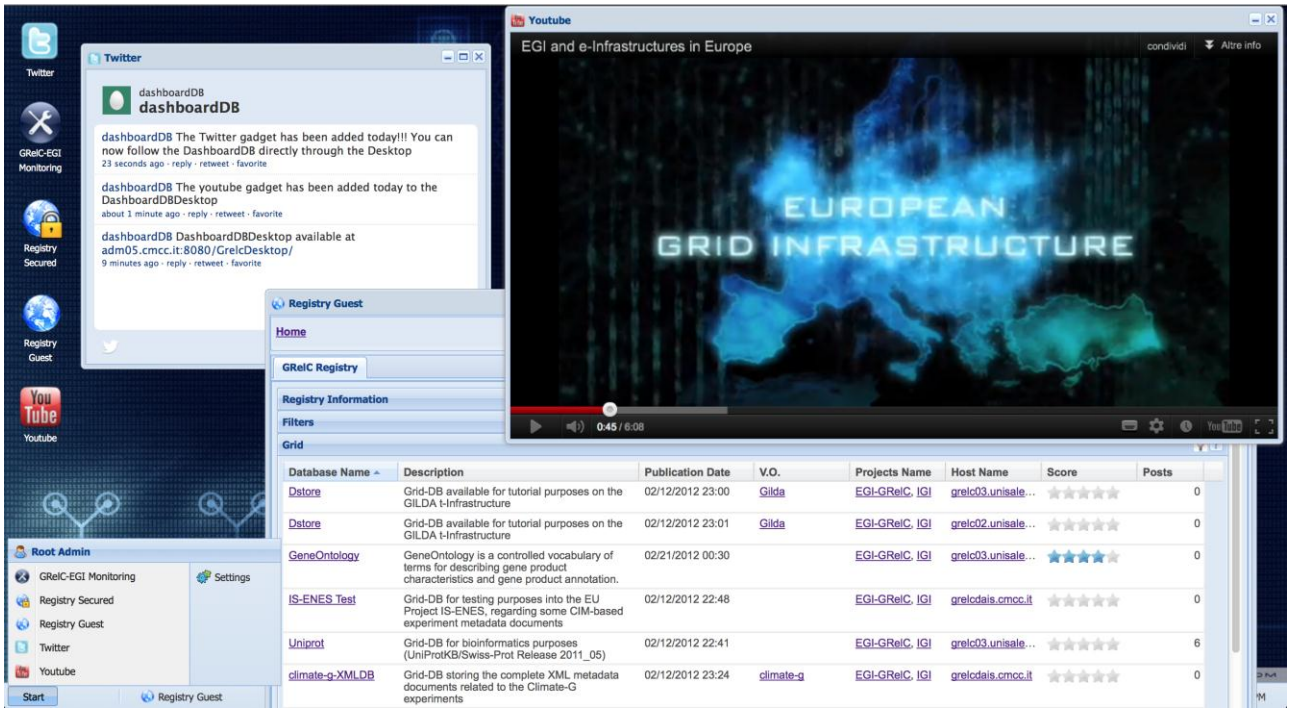

Figure 2: GRelC Desktop environment showing the Registry, Twitter and Youtube gadgets

A second release of the GRelC desktop application has been carried out during PY3. The second release has proved to be a production-level environment.

Concerning point 5 (*dissemination activities*) during PY3, several grid-database services and data providers have been contacted to register/publish their own data resources/services into the DashboardDB system. Some sites (like one in Catania - INFN-CATANIA - and another one in Naples - GRISU-NAPOLI), have respectively updated and installed the latest version of GRelC and registered their service instance on the DashboardDB system. Another dissemination task has been the preparation of a short overview related to the two main GRelC gadgets (DashboardDB Monitoring

and Registry). This document has been prepared during PY3 with NA2 representatives and posted on the EGI website[12].

Another activity related to the dissemination task has been the participation to ICT conferences like SuperComputing 2012 (November 10-16, 2012, Salt Lake City, USA). A video related to the DashboardDB application and the GRelC Desktop has been displayed in the IGI booth. Dissemination activities include also some talks held at the Euro-Mediterranean Centre on Climate Change (CMCC) about the GRelC service and its main application. As a consequence of these meetings, starting from February 2013, the GRelC service is being evaluated: (i) as grid metadata service in the context of a national project on "Situational Sea Awareness" (3-year project started in 2012), and (ii) as grid front-end to a data analytics service for analysis and mining of climate change data (NetCDF format) into another national project (4-year project, started in 2011). A couple of new grid-databases related to multidimensional data have been published in the DashboardDB registry.

Finally, minor dissemination activities are related to updated documentation, news and tweets (DashboardDB Twitter account[13]) on the GRelC website and Desktop.

The current work is related to (i) porting activities of GRelC on EMI1 and EMI2, (ii) dissemination activities to better publicize the results achieved during the PY1-PY3 period and attract new users, (iii) registration of new grid-databases and GRelC service instances into the DashboardDB system and (iv) definition of new use cases related to the HUC. Some of these tasks (like dissemination and HUC support) can be considered long term and will continue even after the PY3 period. In this regard, to sustain these activities, a strong effort in terms of new proposals has been key over the last three years to support - at least over the next two years - the implementation of new application-domain use cases involving the GRelC software (oceanography, fire danger prevention and climate change represent the three most relevant examples).

6.2 SOMA2

SOMA2[14] 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.

The main achievement during PQ9 was a new release of SOMA2 (v1.4.1: Aluminium). In addition to bug fixes, this version added support for user generated proxy certificates in Grid use via SOMA2. This release also contained grid enabled versions of SOMA2 demo program descriptions, which make use of the Open Babel program package. These program descriptions are also taken in use in SOMA2 EGI pilot service which was introduced for users in the EGI Community Forum during March 2012. The latter part of PQ9 saw improvements to common UI elements of SOMA2 but this work slowed down a bit due to the summer holidays. Finally, CSC maintained and operated CSC's SOMA2 services during PQ9.

During the first month of PQ10, work centered around making general improvements to SOMA2. The aim was to stabilize the code for a version release. However, starting from September 2012 CSC has totally used the allocated EGI-InSPIRE SA3 funding. This work is therefore now unfunded and the development effort is focused primarily at the national level. CSC will however support the existing

[12] GRelC gadgets: http://www.egi.eu/services/support/gadgets/grelc/index.html

[13] DashboardDB Twitter account: https://twitter.com/dashboardDB

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

SOMA2 services and it is foreseen that this will also suffice for the needs of the international SOMA2 service (SOMA2 EGI pilot).

During PQ11 a public release of SOMA2 version 1.5.0 Silicon was made available on the SOMA2 web site. This release contains all development efforts of PQ9 and PQ10.

6.3 Workflow & Schedulers

During PQ9, work on Kepler (and the Serpens suite for Kepler) focused on improving the existing scenarios and performing further tests of the existing use cases in order to optimize the workflows. These include the existing fusion workflows (described in the previous report) as well as the astrophysics workflow. In particular work started on the extension of the workflow for astrophysics. New possible templates for fusion workflows were investigated, for example a Grid application (Fafner) connected to a HPC application EUTERPE.

In addition, the papers presenting the results of the work were prepared and presented at a range of conferences, including HPCS and ICCS conference. In addition to the above, some bug fixes and feature requests were applied to the Kepler actors.

During PQ10 work involved the integration of Kepler with GridWay services. This included the development of the actors and workflows for interacting with GridWay using the GridSAM BES interface implementation. Small fixes were made to the Astrophysics workflow in response to user requests and extension to this usecase, developed and reported in previous deliverables, were applied. Furthermore, recent work was described in an article prepared for the Fundamenta Informatica journal publications.

During PQ11 work focused on the integration of Kepler-GridWay. The new OGSA-BES interface to the GridWay meta-scheduler was selected as a mechanism to efficiently delegate the job management. This is justified because the BES interface is an accepted standard, based on the JSDL format, which has demonstrated high performance with GridWay in grid environments and is compatible with different infrastructures. Other solutions require one to install Kepler in the same machine as a GridWay instance, which limits the deployment of the solution. Thus, new Kepler actors have been implemented to bind GridWay BES interface in order to submit jobs, check their status and store IDs, as well as other tools necessary to simplify the creation of JSDL files or the automatic delegation and renovation of user credentials. A new Kepler module has been created, including actors and examples. Additionally, in this period Astrophysics workflows involving the submission of large number of jobs have become possible.

Current work is mainly focused on developing the scenarios using the Kepler-Gridway and on providing the fault tolerance framework around the basic actors and workflows. Furthermore, articles for the FI are being improved as a result of the review process. The poster presenting the work of the integration with Gridway will be presented at the EGI Community Forum and a publication summarizing that is under the preparation. It is also planned to evaluate the solution with scientific use cases scenario.

6.4 MPI and Parallel Computing

The MPI and Parallel Computing 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. These products are also intended to have an impact on other user communities. 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 EGI Community Forum and the EGI Technical Forum, 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;
- Outreach and dissemination at major EGI events and workshops;
- User community, NGI and site engagement, gathering direct input;
- Participation in selected standardisation bodies.

To date, over 120 EGI Sites using the gLite/UMD middleware support MPI.

During PY3 CSIC provided effort to the EGI-INSPIRE MPI Virtual Team. Actions completed included: improved documentation on the EGI wiki, new Nagios probes, Information System correctness improvements and batch system MPI support improvements. A full overview of the activity and actions can be found at https://wiki.egi.eu/wiki/VT_MPI_within_EGI

UNIPG provided MPI support on a best effort, unfunded basis. UNIPG has disseminated information regarding the use of MPI with molecular science applications on the Grid at the ICCSA 2012 Conference in Salvador de Bahia. A report was compiled detailing UNIPG's activities (linear algebra routines, quantum reactive scattering programs and secondary pollutant production Chimere package) and promoting the use of MPI on supercomputers from the Grid. In addition, a white paper is was drafted, outlining the strategies adopted to build a computational chemistry VRC.

TCD helped to establish, and currently leads, the EGI VT-GPGPU virtual team. The team aims to collect detailed requirements from existing and new EGI user communities and their support teams on using GPGPU services in the European Grid Infrastructure. The requirements will used by the EGI Operations community (through the OMB), the EGI User Community (through the UCB) and the EGI Technology Community (through the TCB) to define and implement extensions in the EGI e-infrastructure services in order to meet the communities' demand for GPGPU computing.

During PQ11, the SA3 MPI activity continued with the development of the new probes for MPI. These probes are now ready for deployment in the production infrastructure as soon as the new MPI service type is included in GOCDB. CSIC is monitoring the process in collaboration with partners from the MPI-VT.

CSIC has also created a page in the EGI Wiki that collects and provides information about the MPI services and support mechanisms for application developers and resource providers. This was one of the open actions from the MPI-VT. The page contains links to the relevant documentation for both users and administrators about the MPI services in the EGI infrastructure.

The SA3 MPI team members have continued with the support in GGUS of MPI related issues via the MPI support unit.

7 CONCLUSIONS

During the third year of EGI-InSPIRE, the SA3 work package built on and extended the achievements of PY2. Although domain-specific support is still required, further areas of commonality have been found and the communities understand that shared solutions are much more likely to be supported in the long-term than those that are highly-specific to a given VO. Further advances in this area are therefore possible for PY3 and beyond – some of these are highlighted in the Technical Evolution Group reports of the WLCG project.

Sustainability has been a key concern that has been addressed at the technical and strategic level. Sustainable does not, however, mean self-sustaining: all of the domains supported are dependent on external funding and this is committed in the short-term, planned in the medium and expected even in the (very) long-term. One of the youngest domains – HEP, which dates back some 60 years and is hence just an infant compared with e.g. astronomy – has a vision for the next 40 years, as presented in the CERN Director General's annual address in January 2012. Another positive is the trend towards highly distributed computing and "cloud-like" models is unlikely to be reversed.

The success of achieving a high quality of service with vast amounts of CPU delivered accompanied by hundreds of petabytes of storage and multi-gigabit networking across a heterogeneous federated infrastructure with loosely coupled management domains is now much more than an existence proof or even mere state-of-the-art: it is a reliable, performant and operational sustainable production system capable of meeting the needs of thousands of scientists worldwide 24 hours per day and close to 365 days per year.

8 REFERENCES

MS112	Quarterly Report 5:May 2011 – July 2011 https://documents.egi.eu/document/723
MS113	Quarterly Report 6:August 2011 – October 2011 https://documents.egi.eu/document/881
MS114	Quarterly Report 7:November 2011 – January 2012 https://documents.egi.eu/document/999
MS609	HUC Contact points and the support model: https://documents.egi.eu/document/419
MS610	Services for High Energy Physics: https://documents.egi.eu/document/540
MS611	Services for Life Sciences: https://documents.egi.eu/document/683
MS612	HUC Software Roadmap: https://documents.egi.eu/document/684
MS614	HUC Software Roadmap and Sustainability Plan: https://documents.egi.eu/document/746
D6.3	Annual Report of PY1: https://documents.egi.eu/document/312
D6.4	Capabilities offered by the HUCs to other communities: https://documents.egi.eu/document/472
TEG REPORTS	See http://indico.cern.ch/conferenceDisplay.py?confId=158775.
WLCG 2012	See https://indico.cern.ch/conferenceDisplay.py?confId=146547.