EGI-InSPIRE

SUSTAINABILITY PLANS FOR THE ACTIVITIES OF THE HEAVY USER COMMUNITIES – FIRST ANNUAL REPORT

EU DELIVERABLE: D6.2

Document identifier:	EGI-D6.2-V0.99
Date:	**31/03/2011**
Activity:	**SA3**
Lead Partner:	**CERN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/309

Abstract

The Heavy User Community activities are planning for a sustainable future – either within their own community, as part of the generic production infrastructure, or through some other mechanism. This annual report exposes the sustainability plan for each task within the SA3 activity and assesses the progress made to date towards sustainability and the plans for improving sustainability in 2011.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Jamie Shiers	EGI. eu	21/01/2011
Reviewed by	**Moderator:** Steve Brewer **Reviewers:** John Walsh Wilhelm Buehler	EGI.eu TCD/NA3 NGI-DE/NA3	10/02/2011
Approved by	**AMB & PMB**	-	31/03/2011

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	27 January 2011	Version 0.93 – submitted to review	SA3
2	17 February 2011	Version 0.95 – responded to reviews	SA3
3	21 February 2011	Version 0.96 – clarify that this is an annual report	Jamie Shiers / CERN
4	28 February 2011	Version 0.98 – all responses to reviewers' comments integrated	Elisa Lanciotti / CERN

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:

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.

VIII. EXECUTIVE SUMMARY

This annual report exposes the sustainability plan for each community supported by the SA3 activity. It presents the progress made to date towards sustainability and provides plans for the next year that will improve sustainability. The report covers the activities for High Energy Physics (based on the Services for High Energy Physics described in MS603 [MS603]), for Life Science (based on MS604 [MS604]), Astronomy and Astrophysics, Earth Science and finally the shared services and tools that are used by a variety of communities, including but not limited to those mentioned above.

The primary focus of this document is to expose the strategy for achieving sustainability, rather than detailed work-plans and results. Further details of these are or will be provided in the project annual report for 2010, deliverable D6.1 “Capabilities offered by the HUCs to other communities” [D6.1], as well as earlier project quarterly reports.

Table of Contents

1 INTRODUCTION

The SA3 activity within EGI-InSPIRE provides continued support for activities previously supported by EGEE while they transition to a sustainable support model within their own community or within the production infrastructure by:

- Supporting the tools, services and capabilities required by different Heavy User Communities (HUCs);
- Identifying the tools, services and capabilities currently used by the HUCs that can benefit all user communities and to promote their adoption;
- Migrating the tools, services and capabilities that could benefit all user communities into a sustainable support model as part of the core production infrastructure;
- Establishing a sustainable support model for the tools, services and capabilities that will remain relevant to single HUCs;
- As a further step, seeking additional areas of commonality that would result in lower on-going support costs through reduced diversity.

All of the above mentioned tasks aim at reducing the cost of new developments and support, and have as a common objective improving the sustainability of the model. Working towards sustainability is therefore a key element of the strategy of this activity and this document describes the progress to date in achieving this objective together with plans for the coming year. The document is structured by HUC – a description of the communities involved is given in the section below.

The outlook and technical plans for 2011 will be further elaborated in the HUC Software Roadmap document MS606 [MS606]. Updated versions of this report will be written in 2012 (D6.5) and 2013 (D6.8) due in project months 21 and 33 respectively.

2 HEAVY USER COMMUNITIES

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. These communities focus on domain specific issues, such as how to access High Energy Physics (HEP) applications on EGI, how to enable new physics experiments on EGI and so on.

On the one hand these teams are operated by external projects, such as the Worldwide LHC Computing Grid (WLCG), but on the other hand have members in the SA3/WP6 work package of EGI-InSPIRE. The effort of the distributed WP6 team of EGI-InSPIRE is targeted towards the provision of shared services that will ease the porting of new applications from these scientific domains to the wider grid by detecting and exploiting commonalities between virtual organisations (VOs) and driving the implementations to a generic direction.

At the same time inter-VO collaboration typically not only result in more powerful solutions, but also saves significant amounts of manpower in the long run. Such benefits would be unlikely to be achieved with generic support structures, both for individual large communities such as HEP (which could otherwise develop multiple similar solutions to basically common problems), as well across disciplines (e.g. the usage of Dashboards [DASHBOARD], Ganga [GANGA] and HammerCloud [HAMMERCLOUD] across communities).

In conclusion HUCs can offer benefits not only to new adopters of grid technology but also to each other. This continues to be demonstrated, both by the adoption of tools initially developed for one community spreading to others, as well as at the conceptual level: which offers a more pragmatic solution for existing communities.

2.1 Who are the Heavy User Communities in EGI-InSPIRE?

The purpose of this section is to offer an overview of the current HUCs in EGI and the capabilities and services they offer (see
Table 1). Additional information can also be found in the Work Package Wiki [WIKI].

The communities identified as Heavy Users Communities (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 (CheMIST)
- 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) [MS603]	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. The services managed by the HEP HUC can be classified into (see also table 2): • **Data Management** • **Analysis Tools and Support** • **Persistency & Conditions** • **Monitoring**

Community	Description, capabilities and services offered
Life Sciences (TSA3.4) [MS604]	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 the health-grids. A non-extensive list of the services supported and / or developed by this HUC is given below: • **Grid database management service:** The GRelC service will be exploited to implement some LS specific use cases. • **Metadata catalogues:** AMGA • **Workflow engines:** MOTEUR2 and Taverna (http://www.taverna.org.uk/) • **File encryption service:** Hydra • **Monitoring:** Dashboard and Hudson infrastructure monitoring
Astronomy and Astrophysics (TSA3.5)	Originally, Astronomy and Astrophysics (A&A) institutes usually acquired the necessary resources on local computing facilities and quite often also contracted the access to a pool of resources at supercomputing centres. However, some A&A applications grew in complexity and were considered very suitable to be run on grid infrastructures. At the moment the A&A HUC is devoted to the evaluation of different solutions for the gridification of a rich variety of applications. In particular the work is focused on: • **Visualisation tools:** VisIVO (http://visivo.oact.inaf.it/index.php) for the visualisation of data collected by different A&A projects • **Databases and catalogues:** Currently the HUC is evaluating the possibility of adopting and tailoring different tools already in production (e.g. AMGA, GRelC, etc.). • **Parametric applications:** Use DIANE to study parametric space applications. Additionally, one of the main goals is the accomplishment of a good level of interactivity 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 EGDR service to provide access from the grid to resources within the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR). The ES HUC also includes researchers and scientists working in the climate change 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 – Overview of the Heavy User Communities in EGI-InSPIRE

3 SERVICES FOR HIGH ENERGY PHYSICS

3.1 Introduction

The EGI-InSPIRE WP6-SA3 activity supports or contributes development effort to a wide range of services that are critical to the LHC experiments. Table 2 – updated from MS603 reflects the activities of additional personnel that have recently joined the project – summarizes the services supported by SA3 and their corresponding subtasks. This section gives a brief overview of each of the services supported as part of the Services for High Energy Physics (HEP) task (TSA3.3). It focuses on the objective of moving to a sustainable support model. The overview includes Ganga and the HEP Experiment Dashboard, as used by the HEP community but which are also part of the Shared services and Tools task (TSA3.2).

Service	**Subtask**
Monitoring	
Experiment Dashboard	TSA3.2.1
Analysis tools and support	
Ganga	TSA3.2.2
HammerCloud	TSA3.2.2
AliEn	TSA3.3
CRAB	TSA3.3
CRAB analysis server	TSA3.3
Data management	
ATLAS DDM	TSA3.3
PhEDEx	TSA3.3
AliEn	TSA3.3
DIRAC	TSA3.3
Persistency and conditions	
CORAL	TSA3.3
CORAL server	TSA3.3
POOL	TSA3.3
COOL	TSA3.3
FroNTier/Squid	TSA3.3

Table 2 - Services for HEP supported by SA3

2010 saw the first data-taking run of the LHC between accelerated beams (proton – proton from late March until November followed by several weeks of data taking from lead – lead collisions), following a short run at injection energy in late 2009. All computing aspects required to support these activities have been considered successful and are supported by this activity, along with other partners. Whilst there were a number of service problems, the response to and resolution of such incidents was considered correct, even though some incidents took longer to resolve than the targets.

Given that data taking in 2011 is expected to result in increased load on the grid infrastructure, it will be important to ensure that the services continue with at least the same level of reliability as was seen during 2010.

The challenges of real data taking, processing and analysis have highlighted some gaps and weaknesses that need to be addressed and which will be a priority for the work-plan of the coming year, including addressing scalability and stability issues. In particular, the data volumes and rates in 2010 were significantly higher than those foreseen and the relatively smooth production is a tribute to many years of extensive planning and testing.

It has been recently decided that the LHC will run in both 2011 and 2012, whereas previously a prolonged shutdown had been foreseen in 2012. This decision has an impact on the Services for HEP task as the experiments will be unlikely to wish to deploy large changes during a data-taking period, where the pressure of keeping up with the machine is high. Incremental changes will be acceptable but larger changes that might be disruptive to data processing or analysis would not be.

3.2 General Strategy for Sustainability

As shown in the Table 2, the different services supported for the HEP community through this work package are:

- Monitoring (including, but not limited to the Experiment Dashboards);
- Analysis tools and support (including, but not limited to Ganga and tools built on it);
- Data management (with both WLCG/HEP and experiment-specific components);
- Persistency and conditions (as above).

There are two primary themes that flow through all of these tasks that are of particular relevance to sustainability, namely:

1. All tasks are carried out as collaborative projects, involving people from different sites and institutes as well as the LHC experiments, details of which are provided below;
2. With significantly increased emphasis during EGI-InSPIRE, common solutions – or at least strategies – are sought between the various different tools and the experiments / VOs that are supported. This increased commonality is an important result from the first months of this task.

The lead partner in each service may differ and / or the service may consist of a “common core” plus VO-specific additions or plug-ins.

In all of the services described below, one can expect to see periodic development cycles, as new requirements emerge – e.g. from scale or from new Use Cases – followed by periods where on-going support is the main activity. The different domains are somewhat (but not completely) coupled, leading to broad peaks of activity, measured on the timescale of 1 – 2 years.

A clear conclusion from the results that have been achieved so far is that a sufficiently strong team that is not tied to a given community will give rise to significant savings through the use of common solutions. This is true not only within a specific community, such as the LHC experiments, but also across disciplines, where deployment strategies, operational procedures and/or tools can be effectively exported or shared.

Furthermore, the use of Distributed Computing Infrastructures (DCI) continues to grow, attracting additional applications within already established communities as well as new disciplines that benefit enormously from the work of the existing ones. By using these infrastructures, the communities are typically able to achieve much more and in a shorter time frame. Examples outside the domain of HEP include humanitarian applications where the processing time can be reduced from many months to a few days or weeks.

3.3 Future Directions

Recently several new technologies have appeared which are applicable in the area of distributed computing and can have a big impact on the handling of infrastructure, processes and dependent technologies. This section tries to highlight a few of the current trends and/or technologies that are currently being developed or tried out in pre-production areas.

The area of virtualisation, de-coupling the actual application in a generic environment from the hardware where it runs on, has already been explored for some time and promises to be the major direction for long term sustainability. We therefore plan to explore this technology as first priority. Currently the major Linux distributions are clearly following or even setting this trend, such as the upcoming Red Hat Enterprise 6. Similar to these evolutions also new projects for virtual appliances (CernVM [CERNVM], StratusLab [STRATUSLAB]) appear which allow users to encapsulate applications, decoupling it from the underlying hardware, thus making them easier to manage and distribute. An area with even greater impact will be the continued evolution of multi-/many-core CPUs. Currently applications can handle the provided hardware, but with increasing numbers of CPU cores new ways of job distribution or execution for better usage of these technologies need to be developed. This includes work on job scheduling, such as the "Whole Node" working group that has recently been started within the HEP community but is likely to be of much wider interest. Connected to this area is the usage of General-Purpose computation on Graphics Processing Units (GPGPU) that allows the execution of highly parallelized applications. Additionally, the so-called "NoSQL" database technologies such as Cassandra [CASSANDRA] and CouchDB [COUCHDB] can be employed by services requiring highly scalable datasets.

Finally a field that will have impact on distributed computing is the usage of new file systems and caching mechanisms for handling large datasets or make easier distribution of software (Hadoop [HADOOP], CernVM-FS [CVMFS], FroNTier [FRONTIER]). Besides virtualisation this is probably the most interesting technology that can bring benefit for the long term sustainability since it moves in direction already adopted by industry.

These emerging technologies are again a very good opportunity to propose common solutions for distributed computing. EGI-InSPIRE provides the proper infrastructure, processes and connectivity throughout HEP and even across HUCs to follow up on the latest trends and propose common solutions. Since the beginning of the EGI-InSPIRE project, the sharing of knowledge across different activities has already born fruit. This has led to common solutions in all of the already existing services to which this effort contributes – Monitoring, Analysis, Data Management and Persistency – that would not have happened otherwise. Not only is this a very positive result after relatively few months but also it shows clear promise for the future.

Thus, the project will enable us to:

- Firmly establish information flow between previously disparate tasks in the same area, such as the different distributed analysis tools;
- Identify areas were common solutions can be developed and/or adopted;
- Progressively decrease the maintenance load as a result.

It will not be possible to re-engineer some of the main applications and tools within the duration of the project – principally in the areas of data management and analysis support – and this remains a challenge for possible future projects.

3.4 Project Structure - Benefits

The remainder of Section 3 briefly describes how each of the "projects" (monitoring, analysis, data management, persistency) is organised. This list is not exhaustive but corresponds to those currently supported through EGI-InSPIRE for the HEP community.

In addition to the description of the project organisation, it is important to emphasise how such coordinated support can benefit multiple VOs. A good example that illustrates this is that of "WLCG Operations", which has been in a sustainable state for around two years. Started in conjunction with the WLCG Service Challenges (2005-2006) and continuing through the CCRC'08 and STEP'09 large-scale readiness tests, WLCG operations coordinates the VOs using the "WLCG infrastructure" (which includes inter-operation with non-EGI grids such as those provided by NorduGrid and OSG), together with the sites and services. It has been built up using input from all of these different players and, whilst WLCG continues to play a coordinating role, the bulk of the VO-specific work is now handled within the VOs themselves. Not only does this provide a good example for other areas to follow, but it also clearly shows how working across VOs delivers something more powerful and at the same time more manageable than had each VO developed their own techniques and procedures independently.

3.5 Analysis Tools and Support

Analysis that depends on EGI will continue to be a central activity in LHC computing for the foreseeable future – up to and most likely a few years beyond the LHC machine's lifetime, currently scheduled until 2030.

Even given the expected shutdown of the LHC in 2013, the analysis activity is expected to increase in number of users, number of jobs and amount of data analysed over the next few years. Over this term, development efforts will focus on making more efficient use of the EGI resources to improve analysis throughput and on tuning the user experience in order to decrease the user support load for the tools:

- We plan to evolve the monitoring tools in order to clearly decouple failures originated by the infrastructure services (computing, storage, network) from the ones due to the user application. In will then be possible to rely on a common support model for infrastructure issues while delegating the user specific ones to the single support units inside the VOs.
- To simplify the data access layer we will participate in the ongoing effort for the consolidation of data access protocols. Our contribution will focus on collecting requirements, testing and integration with the existing tools. Reducing the number of protocols will allow simplifying the analysis tools and their external dependencies, with obvious advantages in the long term (think about portability to different Operative Systems).
- In terms of support, we intend to move toward a model of user-to-user support, which will scale positively with the numbers of users (the current centralised support model can hardly be sustained in the long-term since the number of support units will tend to decrease while the number of users will tend to increase). In this respect we plan to provide a framework for user-to-user support, together with documentation and best practices, and coordinate the activity across experiments. A better integration of the experiment specific submission systems with the user support system will be essential.

More in general, concrete steps towards sustainability will focus in improving the scalability, stability and usability of the tools themselves through a collaborative effort between the experiments, key institutes and EGI-InSPIRE SA3 itself. Some of these tools have already been used in production by numerous disciplines outside HEP while others, like HammerCloud, clearly have a wider applicability. Communities that wish to use these tools are expected to participate in the support and maintenance effort. For example Life Science and Earth Science communities are key users and contributors.

3.5.1 Ganga

Ganga has been used for several years for distributed data analysis by the ATLAS and LHCb experiments. There is also a successful record of the use of Ganga by more than 10 user communities from other scientific fields and disciplines and by 500 – 600 users monthly (all communities). The Ganga project leader is located at Imperial College, London whilst development coordinators of the ATLAS and LHCb modules are located at LMU-Munich and Imperial, respectively. Development and user support of the Ganga core is led by CERN IT. The project infrastructure (WWW, SVN, Monitoring, Twiki, test/build machines, etc...) is also provided by the IT department at CERN, with development effort from both IT and the Physics Department. Other developers for the core and experiment plug-ins are located at the Universities of Birmingham, Cambridge, LMU-Munich, Oslo and NIKHEF in Amsterdam. User support for the LHCb and ATLAS communities are provided by and coordinated through the experiments themselves. The long-term sustainability of the Ganga project relies on its modular architecture:

- The Ganga core components include the common software framework shared by all involved user communities and applications, and software process infrastructure (packaging and software distribution tools, release tools and services, internal testing services, central documentation). This commonality ensures that the core packages can be supported with reduced manpower and serve multiple communities.
- Ganga supports several job submission backends, including the most widely supported and utilised job submission services in EGI, i.e. the Cream CE and the gLite WMS (beside Condor and various batch systems). From one side, the support for multiple backends makes the system resilient in the long term for the lack of support of one submission mechanism. From the other side, the agile framework will allow to add with minimal effort new potential backends (a possible integration with Cloud infrastructures is an example).
- Different VOs can integrate Ganga with their native set of services through a plug-in mechanism. This decouples the core components from the experiment specific framework and will be beneficial in the long term in case of important changes in the experiment components, which should be transparent to Ganga core (or at least incorporated with limited effort).

Additionally, the development of Ganga is community-driven with about 5 FTE in total, including those supported by EGI-InSPIRE, spread over some 10 developers around the world. This spread knowledge of the system will mitigate a possible loss of manpower in the future.

3.5.2 HammerCloud

HammerCloud is a distributed analysis stress testing system built around Ganga. Inspired by an older and less advanced service, the CMS Job Robot, it was motivated by a requirement from the ATLAS collaboration for site- and central-managers to easily test a set of grid sites with an arbitrarily large number of real analysis jobs. These tests are useful during site commissioning to validate and tune site configurations, and also during normal site operations to periodically benchmark the site performance.

Development and operation of the HammerCloud service is led by the IT department at CERN (CERN IT) via an EGI-InSPIRE funded position. HammerCloud depends on Ganga for the service delivery as it internally uses the Ganga interface to streamline certain parts of the workflow (such as job definition and interfacing to EGI grid services). Each VO provides operational support for their HammerCloud service instances -- the VO-specific knowledge of these operators is critical to create and maintain useful HammerCloud test templates. The estimated effort per VO is 0.1 to 0.2 FTE. We do not expect these requirements to change over the medium term. HammerCloud has been designed to address real long-term sustainability from the beginning:

- Since it relies of the Ganga submission system, it benefits from the long term support plan of Ganga described above.
- Its modular architecture allows the integration of different analysis framework not based on Ganga. Effort for example has been spent in integration the CMS CRAB system into HammerCloud. The aim is to provide a unique commissioning and site testing framework for all the LHC experiments.

3.5.3 AliEn

The AliEn (ALICE Environment) [ALIEN] framework offers full end-user data analysis capabilities for the ALICE physicists, in addition to the raw data processing and Monte Carlo production management. Its role is firmly established in the ALICE computing model and all large-scale physics analysis since 2006 has been successfully carried out using this framework exclusively. In 2010, the average number of analysis jobs submitted through AliEn and executing in parallel on the Grid was about 3000, in addition to about 2000 jobs of the so called 'Analysis trains', which represent common analysis and data quality assurance tasks run centrally on behalf of the various ALICE Physics Working Groups. This resulted in about 1 job/second completion frequency. The daily average of user analysis jobs during 2010 has been ~180 with ~450 individual users throughout the year. The load from analysis jobs is projected to increase in 2011 and beyond, due mainly to new users and increased data volume.

Even if originally developed for ALICE, AliEn is generic and it can be used by other experiments outside LHC. In particular, PANDA [PANDA] and CBM [CBM] at GSI are also using the same system. Moreover, there have been prototypes of AliEn for non-HEP communities. The core AliEn development team consists of 2 FTEs based at CERN, plus many contributions from participating centres. In 2010, there were contributions from Torino, GSI (Darmstadt), LBL (Berkeley), Armenia and India. The sustainability of the project can be achieved by having the small core team dedicated to coordinate the contributions made by the participating collaborators. On top of that, there are 2 FTEs (one CERN/IT, one from the experiment) dedicated to the operation and coordination of the ALICE sites. Since Oct 2010, the AliEn project leader has been a CERN/IT staff member. Other experiments using AliEn provide their own FTEs for the operation and support.

The plans for improving sustainability of the service foresee some new development in two main directions: introduce a functionality to replicate files automatically depending on their usage, taking advantage of systems already put in place by other VOs; and strengthening the scalability and error recovery of the system. These developments will improve the automation of the overall service, making it more sustainable in terms of support.

3.5.4 CRAB and CRAB Analysis Server

The CMS Remote Analysis Builder (CRAB) [CRAB] was the first analysis tool in CMS to aid users in configuring CMS applications for distributed use [CMSDA], 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 was 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) [CTDR] estimated roughly 100k grid submissions per day. During the second half of 2010 the job submissions routinely exceeded the estimate by 40-50% and CRAB coped well.

CMS has been largely successful in splitting the development and operations teams for CRAB over the last 18 months. The operations team handles the interaction with users, tracks problems and submits tickets, and operates the CRAB servers [CMSAOPS]. The operations teams are generally comprised of contributed effort from participating countries.

The CRAB team has an ambitious program planned in 2011 to release a third generation of CRAB which, besides providing more functionalities and higher reliability, will address several long term sustainability issues:

- The CRAB3 server will share a considerable part of the core code with the CMS WM agent, responsible for the MonteCarlo production. This will include they key aspects of job submission and data access. This will reduce the support load for those components in the long term.
- The CRAB3 server will include a system for asynchronous data movement, based on the gLite File Transfer Server (now supported and actively developed in EMI). This will provide a more reliable handling of output data and therefore easy the operations both from the infrastructure and experiment side (it is in fact one of the major problems at the moment in CMS analysis and requires a lot of operational attention).
- The CRAB3 server will utilize consolidated baseline services for job submission, such as the glite WMS, the CREAM CE, and Condor. Being those services part of several other VOs frameworks (in and outside HEP) portability and support in the future should be better guaranteed (in respect of experiment specific agents and services).

CRAB development is led by CERN IT via an EGI-InSPIRE funded position. FNAL and INFN also contribute ~2.5FTE to the development. The overall ~4FTE are shared by 6 persons in total. The EGI-InSPIRE funded position also contributes to the CMS WM-Core project, which provides the core CRAB functionalities and is also supported by FNAL and Bristol (UK). CERN IT provides the development infrastructure (WWW, SVN, Monitoring, Twiki, test/build machines, etc...). The CRAB server instances are distributed across 4 countries, with half of the server capacity hosted by CERN IT. Operational support for the CRAB server and User support for CRAB are both provided by the CMS experiment.

3.6 Data Management

In 2010 the LHC experiments accumulated some 15PB of data at CERN at rates up to 1.5GB/s per experiment that was distributed across grid sites (giving an additional factor in total data volume) at a total rate that peaked above 4GB/s for one single experiment.

Many of the techniques pioneered and exploited by this community could have significant value in others. Whilst the characteristics of different data widely vary, the need to transfer large data volumes across wide-area networks is already important in numerous other domains. For example, in the area of medical diagnosis and treatment, the latest advances in medical scanning technology ensure the production of higher resolution data, with better quality data.

3.6.1 ATLAS Distributed Data Management (DDM)

The ATLAS Distributed Data Management system is built on top of gLite middleware, managing the experiment's detector, simulated and user data while enforcing the policies defined in the ATLAS Computing Model [ACM]. DDM has been in production since 2004 and the core development team (approximately 4FTEs) is based at CERN with auxiliary effort (approx 2FTEs) coming from affiliated ATLAS institutes. In the last years, one of the FTEs in the core team has been provided by CERN IT and is now funded by EGI-InSPIRE. The main activity of the EGI-InSPIRE funded FTE has been focused on long term sustainability:

- An automated system for site selection and exclusion has been developed. This reduces the amount of operational effort required. In addition, the system is based on consolidated services provided or adopted by EGI such as the CIC portal and GOCDB for the site downtimes, SRM and the BDII for storage space occupancy. Since those services are widely utilised by many communities (including non HEP experiments), the long-term support should be guaranteed. Additionally, several components and collectors have been developed in a general way so that they can be integrated in other VOs frameworks and therefore serve more than one community.
- An automated system for redundant and less utilised data has been developed on top of the SRM based space monitor (described above). Once again, this not only reduces the load on data management operations, but shares core components with other services, to facilitate sustainability in the long term.
- The ATLAS Data Management system has been evolved to rely more heavily on baseline services offered by the EGI infrastructure (such as SRM, FTS and LFC) and supported for several VOs. The amount of ATLAS specific components has been reduced to a minimum and only where really needed. This should facilitate long term support for the core components of the system.
- The ATLAS LCG File Catalog deployment model has been revisited. The various LFCs, scattered around the infrastructure, will be consolidated in a central LFC. This deployment model, already adopted by the LHCb experiment, should be more sustainable in the long term in terms of development (portability), deployment and operations (reducing the effort on various NGIs)

In general, the system has reached a very stable and mature status and is expected to reach maintenance status during the second or third quarter of 2011. The reduction of effort assigned to this complex task will allow the developer to spread the acquired knowledge, work on implementing new solutions for other experiments and develop systems that can be shared by different experiments, as well as communities beyond HEP.

3.6.2 CMS Physics Experiment Data Export (PhEDEx)

The development of the PhEDEx [PHEDEX] system began with the 2004 CMS Data Challenge and therefore precedes the gLite File Transfer Service (FTS), which draws many ideas from it. PhEDEx has been actively developed and has evolved, but the high level architecture deployed in 2004 is still in use today.

The PhEDEx project has been in development for a long time and has successfully transitioned between three lead developers, while maintaining consistency of mission. Currently the PhEDEx development team consists of 2.5FTE spread between EGI-InSPIRE supported effort and that provided by CMS collaboration institutes. The operations team is substantially larger with contributions from CMS central operations, operations from remote sites, and site dedicated local effort to maintain the transfer infrastructure.

The main strong points for PhEDEx sustainability are:

- Decoupling between the core workflow agents and the task agents interacting with the baseline transfer and storage services offered by the EGI infrastructure (such as SRM and FTS). This allows technology-specific transfer agents to developed that are strongly integrated with the baseline services, while at the same time minimizing the impact of technology evolution on the core agents.
- Development of a generic framework to handle local interaction with storage, with plug-ins for storage-specific implementations. Plugins for the most widely-adopted storage solutions in EGI are provided (CASTOR, dCache, DPM, SRM, posix), but it is easy for site-local experts to develop their own optimized plug-ins for interaction with the local storage. This allows community contributions to be easily leveraged when new storage solutions are deployed.
- Development of agents intended to lessen the operational load needed to maintain a local PhEDEx installation, by monitoring agent state, sending automated notifications and taking automatic action such as agent restart. Eventually this could evolve into an intelligent system for automated error correction.
- Deployment of a web data-service to retrieve monitoring information and place new requests into the system. This allows the development of external tools intended to automate operations with little impact on the core components. An example of this is the planned adaptation of the automated system for the management of redundant replicas developed for the ATLAS DDM, which could interact also with PhEDEx through the DataService.
- The agent structure of PhEDEx is extremely powerful, making it very easy to distribute agent-workload over many machines. In the longer term, PhEDEx could go in the direction of factorizing the PhEDEx-specific core behaviour out of the PhEDEx agent, and make it a generic thing. The "agent" model of PhEDEx could be re-used in other applications and other domains.

In general, the system has reached a very stable and mature status and has been proved to scale to data volumes far beyond what is required for LHC data taking. Development effort continues in a few key areas. The 2011 development priorities include a redesign of the monitoring interface to improve the functionality and performance.

3.6.3 ALICE Data Management (AliEn)

AliEn offers several components to manage data movement. The AliEn File Catalogue has been used by ALICE since 2002, when ALICE started using AliEn for the Monte Carlo productions. The AliEn File Catalogue contains all the files registered in the system (more than 200M) and the metadata. On top of the AliEn File Catalogue, there is a full infrastructure to decide the best location to store and retrieve data among the more than fifty storage elements distributed over the world. All of that is done in a transparent way from the user perspective.

As mentioned in 3.5.3, AliEn is also used by several experiments outside the LHC. The people involved in this task are the same as described in 4.5.3. Therefore, the sustainability of the project can be achieved in the same way: a small core team coordinating the contributions from the different collaborators.

3.6.4 LHCb Distributed Infrastructure with Remote Agent Control (DIRAC)

The development of DIRAC [DIRAC] is a shared responsibility among several centres. The involved institutions are: CPPM-IN2P3 (Marseille), Barcelona University, Heidelberg, IHEP (Serpukhov) and CERN. The total effort involved in DIRAC development is about 10 FTEs, of which 1 in the CERN IT department funded via EGI-InSPIRE. Data management is one of the main services for the LHCb experiment provided by the framework. The work for support of the DIRAC data management system in EGI just started at the end of Q2 of 2010. Several tasks have already been agreed with LHCb, the most important ones concerning a new system for consistency checks between the content of the grid storage elements and what is registered in the central file catalogues, and the accounting of the usage of data on the grid storage. Both tasks are crucial for improving the automation of data management operations and the general efficiency in data management.

The development of such systems will first be implemented for the LHCb specific implementation of DIRAC. Subsequently, also other user communities might be interested in their use. For this reason the development is done in the most general way decoupling the settings specific to a particular community from the general functionality. In this way more user communities can take advantage of any new development and there is no duplicated effort to implement the same functionality for more than one community. This approach is highly beneficial also for the sustainability of the project.

While DIRAC was originally developed for LHCb, it is also being used by other HEP communities, including ILC-LCD [ILC] and Belle [BELLE], both with many European institutions participating.

With the experience of extending DIRAC to support other user communities, it became apparent that the DIRAC infrastructure is suitable to support multi disciplinary VOs also outside the HEP domain. This may facilitate the introduction of small communities to the use of the grid infrastructure. Adopting DIRAC is currently under discussion in Earth Science community.

Given the increasing number of DIRAC users, the long-term support of the project becomes of primary importance.

3.7 Persistency and Conditions

Persistency and conditions cover two essential areas for HEP data processing and analysis, namely the handling of the *event data* and *detector conditions data* of the LHC experiments. Event data, that record the signals left in the detectors by the particles generated in the LHC beam collisions, are generally stored on files, while conditions data, that record the experimental conditions (like voltages and temperatures) at the time the event data were collected, are commonly stored using relational database systems. In three of the LHC experiments (ATLAS, CMS and LHCb), some or all of these types of data are stored and accessed inside data processing jobs using one or more of the three software packages developed by the Persistency Framework [PERS-FRAME] project: CORAL [CORAL], COOL [COOL] and POOL [POOL].

The three Persistency Framework packages specialise in different areas of data persistency. COOL, which provides specific software components to handle the time variation and versioning of conditions data, is the only package that is strongly HEP-specific. POOL is a generic hybrid technology store for C++ objects and object collections, using a mixture of streaming and relational technologies.

CORAL, a generic abstraction layer with an SQL-free API for accessing data stored using relational database technologies, is the package that could potentially be most interesting for other communities than HEP (and for this reason will be presented at the EGI User Forum in 2011). CORAL supports several relational database backends and deployment models, including local access to SQLite files, direct client access to Oracle and MySQL servers, as well as read-only access to Oracle through a middle tier server and a caching proxy using the FroNTier/Squid [FRONTIER] and the CORAL server/proxy [CORAL-SERVER] technologies.

The Persistency Framework software has been developed over several years (since 2003 for POOL, since 2004 for CORAL and COOL) through the well established collaboration of developers from the LHC experiments with a team in the CERN IT department (now partly funded by the EGI-InSPIRE project), which has also ensured the overall project coordination. Within the common project, the personnel pledged by ATLAS, CMS and LHCb, coming from a large number of institutes in several countries, have contributed to the development and continue to support the components used by their experiment. The common project, in particular the effort from the IT Department (and from EGI-InSPIRE through it) only deals with components that concern (or, initially, that show the potential to concern) more than one experiment: in the past, individual experiments have already taken up the sole responsibility of components that have proved to be relevant to that single experiment. The usage of the three packages is periodically reviewed with the relevant stakeholders in the IT Department and the experiments; as an example, a new review of the usage of POOL recently started with ATLAS and LHCb to identify a strategy for its longer term support.

While the software is by now mature in its development cycle, a large development and support effort (approximately 4 FTEs) is still required for user support, service operation and maintenance tasks. Software releases are prepared as requested by one or more of the LHC experiments, leading to one release per month on average: this is generally motivated either by functionality enhancements or bug fixes in the Persistency Framework packages, or by upgrades in the versions of the 'external' dependencies (ROOT, Boost, Python, Oracle…). Service operation incidents and user support requests normally result in bug fixes or functional enhancements in the Persistency Framework code, but may often end up in need of a more global analysis involving other software packages (such as Oracle, the grid middleware or ROOT). In particular, understanding service operation issues for CORAL-based applications almost always requires detailed troubleshooting on the Oracle database servers where the data is stored; these tasks are performed by the Persistency Framework team in collaboration with the relevant database administrator team (funded by CERN IT).

The Frontier/Squid system is also relevant to CORAL. Both the software and its deployment are stable and mature, as they have been critical to the operations of the CMS experiment for several years. For CMS, the Frontier/Squid deployment is monitored and supported by a small team composed primarily of institute-contributed effort. The system is now becoming more critical also to the operations of ATLAS, which has recently adopted it to provide data distribution for remote analysis on the grid. This is likely to require additional operational effort, which may eventually be managed in the context of the common project as this activity is relevant to more than one experiment.

3.8 Monitoring

Monitoring of the distributed infrastructure and the activities of the user communities on this infrastructure is a vital condition 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 monitoring of the VOs activities, namely job processing and data transfer. In order to address the first task, a set of functional tests that check the state of the distributed services are performed on a regular basis. The availability of the services and sites is calculated based on the results of these tests.

The Service Availability Monitoring (SAM) [SAM] system which is used for submitting tests, collecting and storing test results, has been re-implemented in order to become compatible with the multi-level monitoring architecture and regional distributed operations model. SAM is the de-facto standard framework service used to monitor EGI infrastructure. The new SAM implementation should improve the sustainability of the system and decrease maintenance and support of the SAM components. The second task, monitoring of the VO computing activities is performed by various monitoring tools which may be coupled with a particular workload management or data management system – for example DIRAC or PhEDEx monitoring. Such systems are described earlier in this document. 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.

3.9 Experiment Dashboard

Experiment Dashboard applications adhere to a set of core development principles: common technology and implementation, loose-coupling to data sources, sharing of monitoring data, and user involvement in the development process. The aims of these principles are to reduce development and maintenance overhead, allow applications to be easily adapted for use by multiple VOs, enable reuse of monitoring data, and ensure that applications meet user requirements. Where possible these principles are embodied in a common architecture and framework. This strategy contributes to the sustainability of the Experiment Dashboard system, while the monitoring functionality of the Experiment Dashboard system contributes to the sustainability of the computing activities of the user communities on the grid in general. The system proved to be an essential component of LHC computing operations and is widely used by the LHC VOs. For example, the Dashboard server of a single LHC VO, like CMS, is accessed by up to 5k unique visitors (unique IP addresses) per month and more than 100k pages are accessed daily. The system covers monitoring of job processing, data transfers and distributed sites and services, measuring their 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. The role of the system is becoming more important with time and this trend is expected to continue for the coming years.

Apart from the development principles mentioned above, the sustainability of the system is ensured by the collaborative nature of the project. The core development and support team of 5 FTEs is located at CERN, one of whom is funded by EGI-InSPIRE. In addition, there is a considerable contribution from the developers of the LHC VOs in particular from ATLAS and CMS, as well as from institutions participating in the LHC project from Russia, Taiwan, India and England. The LHC experiments also take part in operating the Dashboard services, for example as a part of ATLAS computing shift procedures.

The current development directions focus on simplification of Dashboard system deployment and maintenance, further improvement of system performance, and more complete decoupling of the user interfaces from the data repository through the use of new technologies such as AJAX. The latter direction allows the development process to be organized in such a way that it can be performed by several teams independently, with a clear separation of responsibilities. It also simplifies the integration of common user interfaces with various underlying information sources and hence makes possible the adoption of the system by a wider range of user communities.

3.10 Summary of Services for High Energy Physics

All of the services for HEP are currently supported by collaborative projects, which are believed to offer the flexibility and robustness required for sustainability. Manpower remains a critical issue, with resources continuing to shrink at the same time as the LHC machine enters its full exploitation phase from which a vast number of scientific papers and theses can be expected.

As has been shown on many previous occasions, tools and solutions pioneered by this community can be of great benefit to others. Examples in the IT domain include the CERN Program Library, GEANT-4 and most famously the Worldwide Web. The benefit of these to both science and society has been significant – in the case of the Web, revolutionary.

Throughout the remainder of the project, further work on identifying common solutions and techniques will continue, resulting in increased sustainability. It is important to stress that this is enabled through this project and its goals, which provide the necessary impetus and resources.

At a later stage and principally in the areas of data management and analysis, additional rationalisation would be possible, but this is beyond the scope and lifetime of the current project.

4 SERVICES FOR LIFE SCIENCE

Ensuring sustainability in supporting Life Science (LS) grid users' communities is essential to maximize the adherence of the communities to the newly developed infrastructures and allow a more coordinated approach worldwide. To coordinate their efforts and sustain their activity, members from the Life Science community self-organized into the project-independent "Life Sciences Grid Community" (LSGC) over the first period of the EGI-InSPIRE project. The LSGC is currently representing 4 Virtual Organisations (biomed, lsgrid, vlemed and pneumogrid). It receives support from six National Grid Initiatives (NGIs) (Dutch, French, German, Italian, Spanish and Swiss NGIs), the HealthGrid association [HG], and discussions are ongoing with two ESFRI projects (LifeWatch [LW] and ELIXIR [ELIXIR]). A board with representation for each stakeholder has been nominated. The LSGC is coordinated by a leader elected within the community (currently Tristan Glatard from CNRS, France), and scientific sub-area leaders. The LSGC will also be active in supporting the Life Science community through specific VOs, supporting in the negotiations and OLAs with EGI, the NGIs and the grid sites, using the already well-established communication channels provided by the partnering organisations involved in the LSGC. This support will materialize in the establishment of Memoranda of Understanding (MoUs) in the coming months with the grid sites and the NGIs, and also in the development and maintenance of specific tools and services to support the community. Several partners of the LSGC are also partners in the SA3 work package of EGI-InSPIRE.

4.1 LSGC communication channels

Ensuring that information is properly disseminated to all stakeholders helps long-term cohesion of the LSGC. A mailing list of all LSGC partner representatives is used for internal communication. The LSGC has set up a wiki [LSGC] to collect and publish practical and technical information related to the community. A monthly phone conference is organized to address the managerial and technical work involved with representatives from each of these NGIs, VOs and ESFRIs. Minutes of the meetings are available from the wiki. Other means of communications, such as *chat* are also considered as part of the LSGC dashboard development (see Section 7.2 below).

Communication channels will prove essential in maximizing the impact of the tools and services developed by infrastructures such as EGI and the NGIs but also for the communities comprising the LSGC in demonstrating the ported applications. The LSGC in that sense is a privileged environment for dissemination and communication. The HealthGrid association will maintain these channels of communication that will prove useful when negotiating and managing resources with the NGIs.

4.2 LSGC technical team

Efficiently exploiting the grid infrastructure remains challenging for non-specialist end-users such as the ones encountered in the LS communities. An important role of the LSGC to ensure a sustained grid activity is to facilitate access to the grid and improve grid users' experience. The LSGC receives support from the NGIs involved and the HealthGrid association in term of manpower and grid resources. Part of this manpower is used to operate a Technical Team of members from the biomed VO [BiomedTech] to assist the LS user communities. The function of the team is to address problems reported by the community, usually through the GGUS front-line support system. The support is performed using duty shifts. The technical team also anticipates problems by actively probing the most critical services for the proper VO operation through a Nagios server. Procedures have been defined to react to regular maintenance events such as storage element decommissioning operations.

4.3 LSGC user management database

Past experience has shown that managing a VO over a long period of time is a challenging problem. Users are registered in the system upon arrival, but they usually give no notice when they leave.

Their activity might change over time. The LSGC is therefore designing a user management database, which will facilitate liaising with hundreds of users registered in the affiliated Virtual Organisations. The database schema was specified, and the tooling will be developed in the coming months. This user database is considered an important milestone in the setup of a sustainable LSGC. It will interface to a Virtual Organisation Membership Service (VOMS) server as well as the EGI Application Database, to avoid replicating existing information. It will complement the VOMS and application database with extra-information on the users and their affiliations. It will be used to manage the user community and to produce sub-themes mailing lists (per-NGI, per-project, per-scientific domain) to liaise with the end users. It is planned to link information related to:

- User identity (Distinguished Name, email) from the VOMS server
- Application from the application database
- Certificate validity period
- User thematic sub-group
- User NGI
- User grid projects
- All information delivered in the VO registration questionnaire.

Users will register to the “biomed” VO through a single portal collecting all information needed to fill-in the user database. VO managers will be notified of new registrations by email and validate or reject the registration through VO management tools. Upon validation, a registration request will be sent to the VOMS server and the user application description registered in the application database. The user will be granted access to VO sub-groups depending on her thematic sub-group, NGI and projects. A home directory will be created for the user on the LCG File Catalogue (LFC).

The user database will be used to maintain up-to-date list of users and follow-up user activities. The tool will provide list of emails for all the users registered, and sub-lists per thematic sub-group, NGI, and project. It will send yearly notification, upon certificate expiration, for users to update their information and for identifying user leaving the VO. Files belonging to unregistered users should be cleared from grid disk space.

The user database will be interfaced to and monitored from the LSGC community dashboard.

4.4 GRelC data access interface

To address sustainability in the LS context, the use cases involving the Grid Relational Catalogue (GRelC) service will take advantage of general-purpose usage patterns already tested for grid-database management in many domains. A high level of re-usability will drive sustainability, as it has been demonstrated through previous showcases of data and metadata management (see for instance the Climate-G testbed in Section **Error! Reference source not found.** , Services for Earth Science). The SA3 Questionnaire prepared in the GRelC task (Further information on GRelC can be found in section 7.3 below on Shared Services and Tools) will help in the identification of architectural commonalities useful to reduce the time to “gridify” existing LS data sources. The preliminary results from the Questionnaire will be available starting from PQ4.

5 SERVICES FOR ASTRONOMY AND ASTROPHYSICS

Task TSA3.5 is the task dedicated to services for Astronomy & Astrophysics (A&A) HUC. According to the TSA3.5 Work Plan prepared at the beginning of the EGI-InSPIRE project, the activity in TSA3.5 focuses on three main topics:

a) Subtask TSA3.5.2: Visualisation Interface to the Virtual Observatory (VisIVO) and possible other visualisation Grid tools for astronomical data;

b) Subtask TSA3.5.3: Interactivity between Grid-based and HPC-based Distributed Computing Infrastructures (DCIs);

c) Subtask TSA3.5.4: Grid and databases.

The activity planned for task TSA3.5 is primarily targeted at providing these tools and services for the astronomical community; their possible extension to applications of other scientific disciplines, however, is constantly taken into account.

Direct development activity is basically foreseen for subtask TSA3.5.2 only. Activities related to the other subtasks mainly consist in evaluating tools and services already in place and their possible adaptation to A&A applications (this is especially true for what concerns tools and services developed to make grid infrastructures able to interoperate with Databases). In other cases the planned activity mainly consists in providing A&A use-cases and test-beds and in participating in test bed campaigns to prove the effectiveness of solutions adopted in the context of EGI-InSPIRE. This is mainly the case of interactions/interoperability between Grid-based DCIs and DCIs based on other consolidated (HPC) or emerging (Cloud Computing) technologies. Considering the complexity of this matter and the fact that this is one of the key objectives of EGI-InSPIRE at large, task TSA3.5 will collaborate with all other WPs and tasks involved in achieving this goal.

The sections below report on progress in 2010 and plans for 2011.

5.1 *Progress during 2010*

Astronomy & Astrophysics, like Earth Sciences and Life Sciences, is a large community with Institutes and Research Groups in all European Countries; its coordination therefore is quite a challenging task. Thanks to the activity carried out in EGEE projects and now in EGI-InSPIRE, A&A is rightly considered a HUC, able to bring important contributions to EGI and to massively and profitably make use of its infrastructure. Given, however, the intrinsic complexity of this community, both in terms of number and variety of research groups and related research lines, the small participation of the community in EGI-InSPIRE is not sufficient to make the HUC sufficiently strong and well organized as the EGI.eu itself requires to its HUCs. The effort made in EGI-InSPIRE needs to be supported by extra resources and for this reason the second half of 2010 has been spent to identify new funding channels and to prepare proposals to exploit these new opportunities.

Although this activity is not directly related to EGI-InSPIRE, an adequate support to consolidate the results that A&A wants to achieve within this project is of utmost importance.

It is necessary, in fact, to:

a) Strengthen the community in terms of coordination and accomplishment of common, shared objectives;

b) Extend the number of individuals and research groups directly engaged in this main objective beyond the number of people currently involved in EGI-InSPIRE;

c) Ensure the necessary long-term sustainability for all results and objectives achieved by the A&A community with the Grid.

This is indeed what EGI.eu asks to its HUCs; this activity is essential to create the necessary preconditions to make the A&A HUC a well-structured entity able to co-sign formal agreements with EGI. Therefore, any activity aimed at achieving these objectives can be considered a useful support to consolidate the HUC within the European Grid Initiative.

5.1.1 TSA3.5.2: VisIVO and other Visualisation Tools

The Visualisation Interface to the Virtual Observatory (VisIVO) [VISIVO] is a suite of software tools aimed at creating 3D customised views of many type of wide used data sets. Its peculiar characteristic is that there is no limit for what concerns the size of input tables containing data to be processed, thus they are able to support very large scale datasets (tens of Terabytes and more).

The linear scaling observed in our tests suggests that handling large datasets is effectively restricted only by the underlying file system limitations.

VisIVO Server consists of three core components: VisIVO Importer, VisIVO Filter and VisIVO Viewer respectively. To create customised views of 3D renderings from astrophysical data tables, a two-stage process is needed.

First, VisIVO Importer is used to convert user datasets into VisIVO Binary Tables (VBTs). Then, VisIVO Viewer is invoked to display customised views of 3D renderings. VisIVO Filters are collections of data processing modules able to explore datasets enhancing and highlighting their hidden properties.

The software is designed to obtain images and movies from user files. Several data format are currently supported. The process of movie creation could last several hours.

The first fundamental step of the porting of VisIVO in Grid is to obtain, at the end, movies and images directly on the grid storage, even if intermediate files are not produced, and to reduce the overall time for movie production.

For this reason there are two main steps to complete the porting of VisIVO on gLite Grid. First step is to allow the usage of VisIVO directly from a code during the production phase. We are now developing a software layer to use VisIVO directly from codes using the internal arrays without having the intermediate files. A library of VisIVO was designed from November to December 2010 and the development has started.

The second step aims at making possible the executions of the code using MPI, where available. The MPI version of VisIVO is designed for a native parallel cluster, and this means that we need to re-design some parts of the software to complete the porting. However the development of this part must start after the deployment of the VisIVO Library in order to avoid any problem for the new design phase of VisIVO MPI for Grid.

5.1.2 TSA3.5.3: Grid and Supercomputing

During 2010 the current state of the art for what concerns interactions between Grid and Supercomputing and also interactions with the emerging Cloud Computing were evaluated, especially during the preparation of proposals related to new funding opportunities. This evaluation process clearly demonstrated that a lot of work is still necessary in this context, although interesting experiences gained in the framework of past project by other communities with a good chance to recycle the produced results also for A&A. A&A community does not plan to undertake autonomous initiatives in this context. Constant coordination with activities carried out by EGI-InSPIRE at large will continue to provide A&A use-cases and tests-beds which will be used to test implemented tools and services and to propose modifications/improvements if they do not fully satisfy A&A applications. The second half of 2010 has been spent trying to identify significant A&A use-cases and test-beds and in planning a coordinated activity in the context of EGI.eu and of NGIs.

One of the activities that is worth to mention is that carried out by the MPI-WG; members of the A&A community are directly involved in this working group.

5.1.3 TSA3.5.4: Grid and Databases

After the freeze of the development of GDSE (a tool to integrate Databases in Grid proposed by A&A) due to the lack of necessary resources, the A&A community is now evaluating tools and services currently in place to use them in the context of A&A applications. Tools and services currently under evaluation include AMGA, OGSA-DAI, GRelC and others.

5.2 *Activity Plan for 2011*

5.2.1 TSA3.5.2: VisIVO and other Visualisation Tools

During the 2011, the following steps will be accomplished: a) the VisIVO Libraries API (Application Programming Interface) will be completed and fully tested on the gLite infrastructure; b) some use cases will be prepared.

VisIVO Libraries represent a very important step in order to export the usage of the 3D visualisation rendering and movie creation tools to other user communities. When this phase will successfully end, the MPI-compliant version of VisIVO will be tested on the infrastructure, thereafter it will be deployed.

5.2.2 TSA3.5.3: Grid and Supercomputing

According to the activity foreseen for 2011 in the TSA3.5 work plan, the identification of significant A&A use-cases and test-beds for a combined exploitation of Grid and Supercomputing (not within the same applications but rather within the same complex workflows) will go ahead. To this end the involvement of a significant number of A&A research groups (those not directly taking part of EGI-InSPIRE activities) is a necessary precondition. This objective would be greatly eased if the necessary additional resources will become available through the new funding channels identified during the last year. The accomplishment of this activity requires a strong coordination within EGI-InSPIRE given that this is one of the key objectives of the project and is transversal to many of its communities.

5.2.3 TSA3.5.4: Grid and Databases

Integration between Grid and databases remain one of the hot topics for A&A given that astronomical databases are one of the key resources used within the A&A HUC, especially in the context of the Virtual Observatory. Without a good support for databases the Grid does not have chances to penetrate the A&A community. The evaluation of tools and services such as AMGA, GRelC and possibly others will continue. Possible development activities will be agreed with EGI.eu and with NGIs if some of the evaluated tools/services need to be modified to fully meet the requirements of the A&A applications. Also in this case the identification of use-cases and test-beds well representing the majority of A&A applications is a key activity that, in turn, requires a good number of involved A&A institutes and research groups.

6 SERVICES FOR EARTH SCIENCE

The Earth Science task is investing effort into reorganizing and keeping the community in its discipline alive, and to offer services with additional benefit for the community. Earth Science (ES) applications cover various disciplines like seismology, atmospheric chemistry, weather and pollution hydrology (flood and water resource management), ~~a~~nd many others. A critical point for users is the access to existing data infrastructures and centres that have developed many Web services to discover, visualise and download data. The related effort in SA3 is currently centred on the access from the grid to data and resources accessible within the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR). The ES HUC also includes researchers and scientists working in the climate change domain. During EGEE, the Climate-G use case was deployed through several institutes. This use case exploits the GRelC data service for distributed metadata management and the Climate-G portal as scientific gateway for this collaboration. Further information on GRelC can be found in Section 7.3 below on Shared Services and Tools.

Several initiatives and projects have already been submitted to obtain the funding necessary for continuing the community building and productive networking started within EGEE. A common strategy and organisation permits to avoid duplication of efforts, i.e. to prevent the same software model used for various ES disciplines from independently being ported multiple times to the infrastructure. It will permit to explore and develop new services for the benefit of all domains.

There are several primary themes of particular relevance to sustainability, namely:

- Collaboration within the ES VRC
  - Communication channel and contact point
  - Bringing teams working on related topics together to actively support collaboration
- Visibility by the ES community at large
  - Organisation of Grid and HPC sessions within the European Geosciences Union and workshops;
  - Publications and reports highlighting utilisation of the Grid in international journals, such as done in “Earth Science informatics”, two special issues in 2009 and 2010, and “Journal of Hydrology”
  - Flow of information within the community
- Connection with standardisation bodies
  - Outside individual involvements, need to represent interests of ES Grid users in ES and Grid standardisation
- Coordinated technological development to enrich the user experience on Grid, such as advanced data access services and geospatial service integration

Data access and exploitation: Data often represents the result of large amounts of work carried out, which needs to be preserved for traceability and eventual continuation of work, or work based upon these results. The agnostic nature of the distributed infrastructures in regard to specific fields of research and thus data naturally implies that additional services are required to enrich the management of data in these infrastructures with the necessary value-added services.

Geospatial service integration: the ES community at large has actively pursued the goal of shared software and data products for a long time already. One of many important results of this early effort has been the development of standards for geospatial and location based services through the Open Geospatial Consortium. The work goes as far as to allow a completely interoperable service oriented infrastructure based on common web service standards and well-defined interfaces. As a result of this, much expert knowledge and many research products are available today in the nature of geospatial services. Allowing intuitive and efficient usage of these services from the computing infrastructure is beneficial for all users, and important for the acceptance in the community.

The HUC will survey solutions to the above problems and will try to make them available to the users of the EGI if possible. The community will continue the work done in the last year and intends to improve it in a couple of points. Tasks include:

- The web portal (euearthsciencegrid.org) will be extended to include an area for the coordination of work in teams, publicly offering easy tools such as a group workspace and issue tracker for groups working in the area of ES on EGI. This will help in the coordination of work and reduction of redundancy and will hopefully improve the communication between members. This should also provide a basic overview of the projects carried out in the ES Vos.
- The ESR VO for generic Earth Science research will be continued
- Value-added existing software services might be installed for the use in the ES VOs, such as the DIRAC [DIRAC] system, currently under discussion
- Simple demonstration code to discover and download data from the GENESI-DR repository will be published
- Pro-actively contacting projects to leverage common interests and goals
- France Grilles, the French NGI, will support the task through funding effort from CNRS/IPSL to develop solutions to OPeNDAP access from EGI.

7 ADDITIONAL COMMUNITIES & SHARED SERVICES AND TOOLS

7.1 Introduction

This section covers the sustainability plans of those communities not covered by dedicated tasks who are supported by the shared services and tools activity (TSA3.2) of EGI-InSPIRE, as well as for tools that are used by more than one community when the model for sustainability is by tool rather than by community.

7.2 Dashboards

Dashboards[1] are needed to monitor the usage of resource, tools and applications, providing figures and statistics across the board to all stakeholders and users of grid infrastructures and resources centers. Setting up a dashboard is a priority to monitor the use of the grid and other resources by the user communities structuring resources allocations of the resources, negotiations with the resources providers being favored by this transparent way for providing usage information. The Life Science Grid Community is developing a dashboard adapted to the use of its communities and its resources providers with the view of sharing the tools and services with other communities of users such as HEP, Astronomy and Astrophysics, earth observation, biodiversity and eventually ESFRIs in the future. Common features will be made accessible for these communities, the dashboards being developed with a vision of shared services and tools.

- Currently the dashboard is at the phase of design. The analysis of the various actors needing access to the dashboard and the different views and features required to provide these views is ongoing with the communities including the Life Science Grid Community.
- Users: Global view over the state of the grid, number of jobs currently running, status of their own jobs, consumption against their quotas, status of their tickets.
- VO Managers: Status of the VO resources, status of the infrastructure, general usage of the VO resources, statistics per project/users. In addition, the views over the user's management, ticket management.
- Resource sites: Status of their own site according to the LSGC or one of the VOs.

7.3 GRelC

This section provides a brief overview of the GRelC (Grid Relational Catalogue) [GRELC] task and in particular it focuses on the strategy to move towards a sustainable support model by assessing the progress made to date and providing plans for the next years.

Several points discussed below have been, or will be, addressed during the project and represent key factors to really achieve the software sustainability. They are:

(i) General purpose and cross-VO software,
(ii) Collaborative and interdisciplinary approach,
(iii) Community-based patterns,
(iv) From VO to communities,
(v) Robustness,
(vi) High level concepts (communities) and web interfaces (GUI),
(vii) Visibility,
(viii) Sources of funding.

What can be easily argued is that sustainability can be strongly addressed by applying for national and international calls (see point viii) as sources of funding. Anyway, complementary aspects (see

[1] See Section 3.9 Experiment *Dashboard* for the usage and sustainability plans of Dashboards in HEP.

points i, ii, iii, iv, v, vi, and vii) that make the participation and the involvement into new proposals possible are: a larger adoption of the software, a real and effective support in terms of functionalities for HUC issues and use cases, a strong level of dissemination and visibility, an easy integration into existing systems, the scientific value (which can be evaluated in terms of publications on relevant journals), the ability to address common needs across different communities defining some exploitation patterns, the re-usability of the software, its robustness and performance. We do believe that all of these aspects play a significant role and need to be taken into account as well, to address sustainability as it is pointed out in the following sub-sections.

7.3.1 General purpose and cross-VO software

The GRelC [GRELC2] service is a general purpose layer that provides basic functionalities related to the access to relational and non-relational data sources in a grid-based environment. As it can be easily argued, such software is *not tied to a specific community* so it can be exploited to retrieve Astrophysics records, Earth Science metadata, biological sequences in the Life Science context, status-level information from monitoring, accounting and dashboards components, and so forth. This can have beneficial outcomes, such as the definition of common use cases that can act at the same time both as best practices and as effective, tested and well-known exploitation patterns. Such knowledge can be easily exported in other contexts (in a shortened time scale) representing an important point towards sustainability. Furthermore, this *cross-VO property* can attract other users leading to the definition of new use cases across different (and not explored yet) domains and increasing the number of potential interested people. This was the case of the Climate-G testbed (2009), where the Climate Change issues and challenges were (for the first time) addressed at the EGEE level (see Section 6, Services for Earth Science).

7.3.2 Collaborative and interdisciplinary approach

Collaborative projects represent a successful model to address sustainability. Some use cases related to the GRelC service (e.g. Climate-G) have been implemented involving people from different sites/institutions and with different background and expertise. The interdisciplinary benefits of the approach have led to stronger scientific results really addressing users' needs and requirements. This test-bed represents a successful story and it was presented in 2009 during the EGEE-III first year review. These kinds of collaborative projects potentially address sustainability since they can be considered as an incubator for new proposals. Indeed, some of the Climate-G partners were involved into the EU FP7 SAFE proposal (Earth Science VRC) about one year ago.

7.3.3 Community-based patterns

In the recent months one of the goals has been the study of commonalities related to database management across different communities. This has led to an SA3 Questionnaire on database resources that has been sent to the HUCs to understand new and common needs, requirements and use cases. The main goals of this questionnaire are:

(i) To identify common requirements,
(ii) Formalise exploitation patterns,
(iii) Perform a census and define a list of the available EGI databases,
(iv) Define new HUC use cases involving database resources.

Identifying and understanding commonalities in the HUC use cases and reflecting them in new or updated software features is an important result towards sustainability that has been achieved in the first months of this task.

7.3.4 From VO to communities

During the project, the EGI list of databases will not represent a static view about the available data sources (as it will be in the first year inferring the list from the SA3 EGI Questionnaire described above), but a more dynamic, Web2.0 and community-oriented framework instead (in the second and third year). The EGI list of databases (a web-based application named registry) will be searchable and will allow people to post comments and share opinions about the EGI data sources, ask to get access to a specific grid-DB, join a discussion group to exchange feedback, definitively be aware of what other people are doing in the same disciplines, etc. All of these community-oriented functionalities will help both in migrating from *Virtual Organisations* to *Communities* and in answering several questions like:

- Where to post my experience connected with a specific domain?
- How to publicize best practises?
- How to discover new data sources, software?
- Where to find what other people are doing in the same domain?
- Are there virtual places accessible by non-grid users and discussing about or addressing domain-specific issues and challenges?

The planned activities will play a crucial role in addressing sustainability through a community-based support.

7.3.5 Robustness

Making the GRelC software more robust has been an important activity of the first months of this task. Instead of producing new GRelC releases with new functionalities, we basically tested and updated the existing software improving its stability and robustness. Stable and robust software is an important requirement to address sustainability.

7.3.6 High level concepts (communities) and web interfaces (GUI)

Two additional important aspects to address sustainability are the ability to manage communities (which are different from VOs) and to provide high level web interfaces to effectively support them. The activity that has been carried out so far is trying to address both, through a web-based registry (as described in the DoW) that will represent the entry point to the EGI Database of databases. Communities (as they are intended in our task) are not VOs from several points of view: the scale (a community is bigger than a VO), the abstraction level (a community is a higher level concept), the security concerns (a community-user does not exploit proxies or grid certificates), the grid expertise (a community-user could not be a grid user), the background (a community user is not a grid or IT expert). The sustainability of a middleware service depends on the usability, quality, look and feel, completeness, portability and pervasiveness of the associated high level interfaces. A high quality GUI is able to promote (attracting new users) and sustain a middleware level service as needed. This is the case of the of the grid-database registry, part of the activity related to the GRelC service, which is a web application exploiting the Web2.0 principles (e.g. 'mashup', tagging, blogging, etc.).

7.3.7 Visibility

Visibility is another crucial point towards sustainability. It can be addressed through website, wiki, oral presentations, tutorials, publications, organisation of sessions and workshops in conferences. All of these factors address sustainability since they contribute to have a larger user community adopting, exploiting, testing and validating the software.

In this regard the GRelC software has a project website (www.grelc.unile.it) and a wiki on GILDA with several tutorials (https://grid.ct.infn.it/twiki/bin/view/GILDA/GRelCProject). Moreover, in the last years, several training events have been held to attract new users (for instance in the ICITST 2009, GPC09, PDCN2009 and EGEE08 conferences). Several papers related to GRelC have been published in the last years discussing interesting use cases in the Life Science and the Earth Science contexts [GRELC, GRELC2].

Presentations in community-oriented conferences can be also effective to reach this objective. In this regard, the GRelC service has been often presented in Earth Science and Environmental conferences like the European Geosciences Union (EGU) conference in Vienna both in 2009 and 2010. An oral and a poster contributions have been accepted this year too (April, 2011) and will be presented in the **"**Earth science on distributed infrastructures and in HPC" session (in the context of the "Earth & Space Science Informatics" (ESSI) Division) which has been co-organized jointly with some EGI Earth Science representatives.

7.3.8 Sources of funding

Sustainability can be addressed by participating to national and international calls. New proposals can provide the needed funding resources to improve the software, extend it and provide a stronger support to the end-users. The past and current experience with the Earth Science community represents a concrete example. Indeed, there is a strong interaction and expertise in this domain and we expect to participate with this community into new proposals as we did with the SAFE one.

We expect to do the same in the next year(s) also with the Life Science community; right now we are already supporting the LSGC project (see Section 4) providing a dedicated GRelC service instance to implement some database-oriented use cases – see Section 4.4).

During 2010, the GRelC software has been supported by the Euro-Mediterranean Centre for Climate Change (CMCC) project, funded by four Italian Ministries. The support includes the implementation of specific use cases involving the GRelC service to manage distributed metadata for the Environmental community in a grid environment. The follow-up of this project is a grant already approved in 2010 that will start from the second half of 2011 (three years term) and will concern the extension of the GRelC software to include On Line Analytical Processing functionalities for climate change datasets. Moreover, during 2010, the GRelC software has been and will be supported by the EU FP7 IS-ENES project with a small task regarding the management of XML metadata in grid-based environment devoted to climate change.

Finally, during 2010, at a national level we submitted two proposals related to the Call "Programma Operativo Nazionale "Ricerca e Competitività" 2007-2013" and funded by the "Italian Ministry for Education, University and Research" and by the "Fondo Europeo di Sviluppo Regionale" of the European Union, concerning the GRelC service. A response related to these two proposals should be available in the next weeks. With this regard, the plan for the year 2011 is to apply to new PON calls, at a national level, that have been already announced by the Italian Ministry for Education, University and Research and are expected to be published in the near months.

7.4 The CheMIST Heavy User Community

7.4.1 The composition of the HUC

After the negative shock caused by the rejection of the ROSCOE proposal (of which the Chemistry and Molecular Science and Technology (CMST) community was part) to establish a Specialized Support Center for the aggregation of various communities, some best effort (voluntary unfunded) activities have been carried out to ensure continuity to the work of the VOs of the former Computational Chemistry (CC) cluster of EGEE and to the functioning of the hardware patrimony, the availability of the program library, the support to the users. At present a second phase has been started in an attempt to build a new stronger aggregation of the members of the CMST community. The current plan aims to aggregate together the COMPCHEM and GAUSSIAN VOs with other VOs like ENEA-Grid and MoSGrid. ENEA is the Italian Agency for New Technology, Energy and Sustainable Economic Development (http://www.enea.it) operating in the fields of energy, environment and new technologies to support competitiveness and sustainable development having distributed its computational resources among 6 computer centres connected to the public network (called ENEA-GRID) that is interconnected with other Italian and European grid infrastructures. MoSGrid (MOlecular Simulation Grid) is part of the German NGI (D-Grid) Infrastructure operating on Unicore and consists of more than 100 potential users and supporters. This joint endeavour has led to the presentation of a proposal for the November 2010 INFRA call for e-Science environments aimed at assembling a HUC named CheMIST with which other VOs, user clusters and international organisations have expressed their wish to cooperate. Among them members of the Computational Chemistry Division of the European Chemical Society EUCHEMS, of the COST Actions D37 "Grid computing in Chemistry: Gridchem", CM 702 "Chemistry with ultra-short pulses and Free-Electron Lasers: Looking for Control Strategies Through "Exact" Computations", CM901 "Detailed Chemical Models for cleaner combustion", of the European Chemistry Thematic Network (ECTN) and of the International Consortium of the Erasmus Mundus Master in Theoretical Chemistry and Computational Modelling (TCCM).

7.4.2 Hardware and Infrastructure

Hardware and infrastructure available to the community are at present those committed by the partners to the European Grid Initiative and the fraction of the Grid platform made available to the involved VOs and user clusters. A development of the MoU between Heavy User Communities and EGI will be concerned also with the nature of the hardware and infrastructure made available by the HUC.

7.4.3 Utilization and Improvement of Existing Grid Tools

Among the basic Grid tools developed by other HUCs to the end of supporting their users, to develop workflows, to manage repositories the intra-VO tools: P-GRADE, MPI, Ganga as well as the more general grid tools DPM, LFC, FTS, Nagios are being regularly used. At the same time other tools (such as Kepler, and SOMA2) are being considered for future usage while the use of some others devoted to monitoring such as Dashboards, VisIVO, Ganga, HammerCloud, and CRAB, Data management (ATLAS, CMS PhEDEx, AliEn, LHCb DIRAC,…), Persistency and conditions (CORAL, COOL, POOL, Frontier, …) have not been considered.

In the meantime, CYFRONET is continuing to work, on a best effort basis, on InsilicoLab, a work environment for chemists allowing to plan, create, modify and execute complex scientific experiments in chemistry. Along this line a new search tool has been recently added to better support accessing files in LFC replicas. Several bugs have been fixed especially those related to Turbomole execution. A new version of CCLIB library has been added that greatly increases the extraction of data from Turbomole output files. A work has been started on enabling parallel execution of all packages offered via Environment. In parallel basic integration of the Gromacs package has been started.

In addition to providing MPI support within the EGI-InSPIRE, UNIPG offers improved support, on a best effort basis, to its growing community. This is provided through the development of its embryo Grid Framework, GriF, devoted to the assistance of the users in Grid resource selection for the optimal use of the Grid platform and the evaluation of quality of Grid services.

University of Cologne also offers to its community the WS-PGrade UNICORE Submitter and Portlets as graphical user interfaces to software packages supported by the Project Molecular Simulation Markup Language for the description and results of molecular simulations. It also offers the design and implementation of data repositories storing and providing single computational tasks, workflows, and results of simulations.

ENEA offers to its users the virtual laboratory CMAST (Computational MAterials Science and Technology, www.afs.enea.it/project/cmast). CMAST provides a virtual environment where researchers, from both Universities and industry, work together by sharing skills, software, specialised services and best practices. CMAST supports its partners in efficiently planning and managing large-scale simulations using the ENEA-GRID.

7.4.4 Design and Implementation of New Grid tools

Other software of specific interest to CheMIST HUC will be either designed from scratch or developed from experimental tools for the specific goals of:

- Providing the services necessary for the assemblage and the internal and external utilization of multi-scale and multi-physics large applications (grid empowered molecular simulators, atmospheric secondary pollutants production prediction, clean combustion modelling, chemical products toxicity evaluation, Chemistry learning objects repository and self-evaluation procedures);
- Supporting the research necessary to define and develop interoperable data models and building blocks (programs and procedures) of the computational engines needed to empower efficient molecular simulations aimed at producing scientific and technological innovation in the CheMIST knowledge area;
- Via the definition of a Quality of User (QoU) and Quality of Services (QoS) based credit system properly rewarding the various activities (research, production, dissemination, etc.) and assigning resources.

Along this line the adoption of FARO (the ENEA portal) and KEPLER (the fusion workflow) will be generalized and a comparison of related performances with other instruments will be carried out.

7.5 Fusion

7.5.1 Kepler/GridWay

Kepler is open source middleware covered by the BSD licence. The Kepler software is developed and maintained by the cross-project Kepler collaboration, which is led by a team consisting of several of the key institutions that originated the project: UC Davis, UC Santa Barbara, and UC San Diego. There is a large open community around Kepler, using and extending Kepler since many years. Also a number of collaborative projects support Kepler development, which is important to offer the robustness required for sustainability.

There are two main stable, self contained and easily extensible releases of Kepler. Both versions can be used as a base for own workflow development or extension base for any community.

Serpens [Serpens] is an open source module for Kepler initially developed in the EU Euforia project, that handles interactions with different grid middlewares. Serpens provides templates for different types of workflows scenario. Serpens relies on underlying grid middleware, so it will have to follow changes in the grid middleware. Usage of the standards is believed to give certain level of sustainability. However minimal manpower will be required to maintain the changes. It is believed that the effort required for this will be distributed among the fusion (and possible other) community. Considering this contribution is not specific only to the fusion community, other communities will be able to take advantage of this task.

Also the module that enables the interaction with GridWay will require in the future minimal manpower in order to track the changes in the underlying middleware. GridWay is an open source metascheduler included in the Globus Toolkit also able to interact with gLite-based grid infrastructures. GridWay provides a command line interface and also DRMAA [DRMAA], an OGF standard for job management.

It is assumed that the effort required for this interaction will be distributed among the fusion (and possible other) community. The interaction between Kepler and GridWay should be long-term sustainable since it relies on standards such as DRMAA. In case of changes, little manpower will be required to incorporate the modification to the developments.

Kepler is de-facto standard for the European and US Fusion communities. Kepler workflows are planned to be used in this community for longer time scale. This assures the possible usage of this integration in this community. The usage of these tools will be devoted to connect codes that act on different time and space scales as well as codes that deal with different physics phenomena. This latter target is basic for the task of building the numerical fusion reactor, which will allow one to predict the performance and behaviour of these future devices.

EGI-InSPIRE provides a mechanism for sharing the tool and facilitating future development within different communities. Already there were interest and technical discussions with the Chemistry community about usage of Kepler for some of their applications. There are plans to continue such collaboration and search also the possibilities for further collaboration with other communities.

Kepler is an open source project and has a large community built around it. This guarantees the Kepler sustainability. Serpens module is maintained by fusion community and the development is not limited to the scope of EGI-InSPIRE project. Serpens gathered some communities around the tool and is supported by some projects and volunteers from different institutions. It is anticipated that during the EGI-InSPIRE project, the tool will be adopted by other communities who will join the Serpens and Kepler communities and subsequently support and use these tools.

7.6 MPI

The MPI sub-task focuses on a number of objectives over a 36-month period. This sub-task is producing numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry (see Section 7.4 above), 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 User 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.

UNIPG, CSIC, INAF and TCD partners have a great wealth of experience in designing, producing and deploying MPI applications under gLite. These range from relatively simple codes, to large-scale production workflows using multiple externally provided (and widely used) MPI-enabled libraries. TCD and CSIC will also engage with the ARC and UNICORE communities, and will produce high-level documentation for MPI application development and submission under these middleware. In project year one, the first MPI “cookbook” will be produced addressing MPI application development and MPI job submission. This shall be reviewed and updated during project years two and three, as expected middleware changes and new features, such as generic “Parallel” application support, OpenMP support, and Graphics Processing Unit (GPU) application support are introduced. In the past year, there has been a wide-scale increase in the number of sites supporting GPU clusters. There is currently no native support for GPU resource integration in the middleware stacks, in particular, at the glue schema (information system) layer and at the batch system layer. As a result, in order to avoid repeating mistakes made in the early years of MPI integration, effort will be required to ensure that the middleware and sites support their integration in a consistent, uniform manner. Both interaction with the middleware developers and the availability of the MPI cookbook will build awareness and expertise in grid-enabled MPI –an important requirement for distributed systems. CSIC and TCD will also proactively engage with related EMI middleware teams and the WLCG “Whole Node” working group.

As part of the User Community engagement effort, the MPI team will regularly survey VOs, Users and Site administrators for critical feedback. This will also act as a means to gather information about current deficits and future requirements. The first survey will be produced in project month 16.

The MPI sub-task will run a training event at the EGI User Forum 2011, and will also present other

MPI related material. MPI training will allow user communities to build up expertise so that the communities themselves can take responsibility for future training.

8 CONCLUSIONS

We have described above the on-going work and plans on moving to a sustainable future for the Heavy User Communities that are supported by EGI-InSPIRE SA3. Whilst it is clear that many of these activities will be required well beyond the lifetime of this project – nominally up to 2030 and beyond in the case of the LHC – good progress has been made in this direction.

On the other hand, not only are some tools used by multiple communities – with their use (slowly) increasing – there are also clear benefits that can result from bringing different disciplines together, as well as offering solutions that can be readily reused by new communities. This is the strongest argument for future common projects in this area, which by construction encourage reuse in a way that vertical funding cannot.

9 REFERENCES

MS603	MS603 – Services for High Energy Physics: https://documents.egi.eu/document/160
MS604	MS604 – Services for the Life Science Community: https://documents.egi.eu/document/236
MS606	MS606 – HUC Software Roadmap: https://documents.egi.eu/public/ShowDocument?docid=310
D6.1	D6.1 – Capabilities Offered by the Heavy User Communities: https://documents.egi.eu/public/ShowDocument?docid=154
CERNVM	A practical approach to virtualization in HEP; P. Buncic et al, to appear in The European Physical Journal Plus
STRATUSLAB	http://stratuslab.eu/
CASSANDRA	http://cassandra.apache.org/
COUCHDB	http://couchdb.apache.org/
CVMFS	Distributing LHC application software and conditions databases using the CernVM file system; J. Blomer et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
HADOOP	http://hadoop.apache.org/
FRONTIER	Greatly improved cache update times for conditions data with FroNTier/Squid; D. Dykstra and L. Lueking, Computing in High Energy and Nuclear Physics (CHEP'09), Prague, Czech Republic (2009)
DASHBOARD	Experiment Dashboard for Monitoring of the LHC Distributed Computing Systems; J. Andreeva et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
GANGA	Ganga: a tool for computational-task management and easy access to Grid resources; Computer Physics Communications, Volume 180, Issue 11, (2009)
HAMMERCLOUD	HammerCloud: A Stress Testing System for Distributed Analysis; Daniel C. van der Ster et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
WIKI	https://wiki.egi.eu/wiki/WP6:_Services_for_the_Heavy_User_Community
CRAB	CRAB: A CMS Application for Distributed Analysis, Nuclear Science, IEEE Transactions on Volume 56, Issue, Part 2, Oct. 2009, pp:2850 - 2858
CMSDA	Distributed Analysis in CMS; Journal of Grid Computing, Vol.8, Number 2, 159-179, (2010) doi: 10.1007/s10723-010-9152-1
CTDR	CMS Computing: Technical Design Report; M. Della Negra et al, CERN-LHCC-2005-023

CMSAOPS	CMS distributed analysis infrastructure and operations: experience with the first LHC data; E. Vaandering et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
ACM	The ATLAS Computing Model; R. Jones and D. Barberis, CERN-LHCC-2004-037-G-085
PHEDEX	Software agents in data and workflow management; T. Barrass et al, Computing in High Energy and Nuclear Physics (CHEP'04), Interlaken, Switzerland, 2004
ALIEN	AliEn: ALICE environment on the GRID; S. Bagnasco et al, Journal of Physics: Conference Series 119 062012 (2008) doi: 10.1088/1742-6596/119/6/062012
PANDA	http://www-panda.gsi.de/
CBM	CBM experiment at FAIR: Acta Phys.Pol. B41(2010)341
MONALISA	MonALISA: An Agent Based, Dynamic Service System to Monitor, Control and Optimize Distributed Systems; I. Legrand et al, Computer Physics Communications, Volume 180, Issue 12, 2009, Pages 2472-2498
DIRAC	DIRAC – The Distributed MC Production and Analysis for LHCb; A. Tsaregorodstev et al, Computing in High Energy and Nuclear Physics (CHEP'04), Interlaken, Switzerland (2004)
ILC	http://www.linearcollider.org/
BELLE	http://belle2.kek.jp/
PERS-FRAME	LCG Persistency Framework (POOL. CORAL, COOL): status and outlook; A. Valassi et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
CORAL	CORAL, a software system for vendor-neutral access to relational databases; I. Papadopoulos et al, Computing in High Energy and Nuclear Physics (CHEP'06), Mumbai, India (2006)
COOL	COOL, LCG conditions database for the LHC experiments; A. Valassi et al, Proc. NSS 2008, Dresden, Germany (2008)
POOL	POOL developments for object persistency into relational databases; G. Govi et al, Computing in High Energy and Nuclear Physics (CHEP'06), Mumbai, India (2006)
CORAL-SERVER	CORAL server and CORAL server proxy: scalable access to relational databases from CORAL applications; A. Valassi et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
SAM	Evolution of SAM in an enhanced model for monitoring WLCG services; Collados D et al, Computing in High Energy and Nuclear Physics (CHEP'09), Prague, Czech Republic (2009)
HG	HealthGrid association, http://www.healthgrid.org
LSGC	LSGC wiki. http://wiki.healthgrid.org/LSVRC:Index

LW	LifeWatch ESFRI. http://www.lifewatch.eu/.
ELIXIR	ELIXIR ESFRI. http://www.elixir-europe.org/.
BiomedTech	Biomed technical team. http://wiki.healthgrid.org/Biomed-Shifts:Index
VISIVO	VisIVO: Data Exploration of Complex Data; G. Caniglia et al, Mem. S.A.It. Vol. 80,441 (2009)
GRELC	The GRelC Project from 2001 to 2011, ten years working on Grid-DBMSs; S. Fiore et al, to appear in book of "Grid and Cloud Database Management", Springer, 2011.
GRELC2	The data access layer in the GRelC system architecture; S. Fiore et al, Future Generation Comp. Syst. 27(3): 334-340 (2011). http://dx.doi.org/10.1016/j.future.2010.07.006.
CMAST	Computational MAterials Science and Technology Virtual ENEA Lab http://www.afs.enea.it/project/cmast
FARO	Fast Access to Remote Objects, an NX based customisable interface to the ENEA Grid resources and services http://www.garr.it/eventiGARR/conf10/docs/rocchi-abs-conf10.pdf http://www.garr.it/eventiGARR/conf10/docs/rocchi-pres-conf10.pdf http://indico.cern.ch/contributionDisplay.py?sessionId=27&contribId=18&confId=69338
Serpens	http://serpens.psnc.pl
DRMAA	Distributed Resource Management Application API http://www.drmaa.org/