EGI-InSPIRE

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

EU DELIVERABLE: D6.3

Document identifier:	EGI-D6.3-V0.28.doc
Date:	**14/06/2011**
Activity:	**SA3**
Lead Partner:	**CERN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/312

Abstract

This public report describes the current status of the tools and services provided within SA3, progress and achievements since the start of the project in May 2010, as well as future activities including measures to assure sustainability of the current work.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Jamie Shiers	CERN/SA3	23/02/2011
Reviewed by	Moderator: Mario David Reviewers: Mike Jones	SA 1.3 NGS	30/5/2011
Approved by	AMB & PMB		14/6/2011

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	23/02/2011	TOC and HEP section	IT-ES group, CERN
2	25/02/2011	GRelC + LS contributions added	Sandro Fiore, Giovanni Aloisio (GRelC), Johan Montagnat, Samuel Keuchkerian
3	03/03/2010	Fusion contribution added	Antonio Gómez, Marcin Plociennik, Paco Castejon
4	07/03/2011	Chemist added	Antonio Laganà
5	16/03/2011	MPI added	John Walsh
6	28/03/2011	ES added	André Gemünd
7	28/03/2011	SOMA2 added	Tapani Kinnunen
8	30/05/2011	A&A. Contents complete.	Claudio Vuerli
9	12/06/2011	Final version	Jamie Shiers

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

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

VI. TERMINOLOGY

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

VII. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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 a 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 document is the first Annual Report on the Tools and Services of the Heavy User Communities.

It describes the activities and results of these communities, together with the outlook for the remainder of 2011. Work on common solutions, together with plans for 2011 are identified.

Additional information, positioning this work in the context of EGI-InSPIRE as a whole, can be found in the project quarterly reports, MS105 [MS105], MS106 [MS106] and MS107 [MS107] for project quarters 1 – 3 respectively.

Table of Contents

1 INTRODUCTION

This activity within EGI-InSPIRE provides continued support for selected 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 EGI 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.

Working towards sustainability is therefore a key element of the strategy of this activity and progress to date in achieving this objective is covered in detail in D6.2 "Sustainability Plans for the HUCs" [D6.2].

The present document is structured by HUC and a brief description of the communities involved is given in the section below.

The work package primarily provides support for the largest communities – High Energy Physics, Life Science, Earth Science and Astronomy and Astrophysics – together with shared tools and services that are also used by others such as Fusion and the CheMIST community. However, the project has already seen increasing interest in a number of tools from outside the discipline that developed it. This, coupled to a significant focus on common solutions within the larger communities, is an important tangible result of the first year of the project.

This document – the first Annual Report on the Tools and Services of the Heavy User Communities – describes the current status of the tools and services provided within SA3, progress and achievements since the start of the project in May 2010, as well as future activities including measures to assure sustainability of the current work.

2 HEAVY USER COMMUNITIES

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

The communities identified as HUCs within this work package are:

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

Community	Description, capabilities and services offered
All HUCs (TSA3.2)	This task provides support for tools and services that are used by more than one community, including Dashboards, applications such as Ganga, Services such as HYDRA and GRelC, Workflows and Schedulers (SOMA2, Kepler, Taverna) and MPI.
High Energy Physics (TSA3.3) [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. Increasing focus is placed on common tools and solutions across these four large communities together with their re-use by other HEP experiments as well as numerous different disciplines and projects.
Life Sciences (TSA3.4) [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.
Astronomy and Astrophysics (TSA3.5)	The A&A HUC is devoted to the evaluation of different solutions for the gridification of a rich variety of applications, as well as the accomplishment of a good level of interoperability among different technologies related to supercomputing, i.e. High Performance Computing and High Throughput Computing, Grid and Cloud.
Earth Sciences (TSA3.6)	Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting and many others. Their presence in SA3 is currently centred in the implementation, deployment and maintenance of the 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 includes also 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 - Summary of the Heavy User Communities

2.1 Collaboration with NA3

The experience that has been gained in supporting the Heavy User Communities throughout the lifetime of the EGEE series of projects would clearly be beneficial to new communities, such as those supported by the NA3 work-package. This is made explicit in the Description of Work for the project, which lists as an objective *"O4: Interfaces that expand access to new user communities including new potential heavy users of the infrastructure from the ESFRI projects"* and proposes:

- SA3 – Dissemination of heavy-user community activity to new users;
- NA3 – User-support mechanisms for Virtual Research Communities (VRCs – expected to be established within ESFRI projects).

The milestones and deliverables that have been completed so far by this work-package directly address these goals (see the following section for a more detailed description) and SA3 and NA3 have sought – within the limits of the available manpower – to collaborate closely on user and community support issues.

The work-package has also been involved – through its constituent communities – in establishing Virtual Research Communities (VRCs), which includes setting up of formal channels between the communities – by definition well established – and EGI-InSPIRE bodies.

The main HUCs are also involved in the nascent EGI-InSPIRE boards, such as the User Services Advisory Group (USAG) and User Community Board (UCB).

A concrete example of focussed collaboration has been around the development of a potential generic dashboard for new user communities, discussed in more detail in section 3.6.2 below below.

2.2 Summary of Milestones and Deliverables

The following table lists the milestones and deliverables that have been achieved so far. Several of these are revised on an annual or semi-annual basis. References to the corresponding report can be found in the table at the end of this document or using the hyperlinks in the online version. Whilst this deliverable is written so as to be largely self-contained, the kind reader is referred to these other documents for supplementary information.

Milestone / Deliverable	Due Date	Lead Partner (#)	Title
MS601	PM1	CSC (13)	HUC Contact points and the support model
MS602	PM4	INFN (21)	HUC Software Roadmap
MS603	PM4	CERN (35)	Services for High Energy Physics
D6.1	PM4	CERN (35)	Capabilities offered by the HUCs to other communities
MS604	PM4	CNRS (14)	Services for the Life Science Community
MS605	PM8	TCD (19)	Training and dissemination event
D6.2	PM9	CERN (35)	Sustainability plans for the HUC activities
MS606	PM10	INFN (21)	HUC Software Roadmap
D6.3	PM11	CERN (35)	Annual Report on the Tools and Services of the HUCs
MS607	PM12	CNRS (14)	Hydra service deployment
MS608	PM12	INFN (21)	Integration of the VisIVO server with the production infrastructure

Table 2 - Summary of Milestones and Deliverables

3 SERVICES FOR HIGH ENERGY PHYSICS

3.1 *Introduction*

Since the start of the project, the work of this task has been predominantly for the 4 main LHC experiments – ALICE, ATLAS, CMS and LHCb. This has been an extremely exciting time with the start-up of the first major new particle collider since over a decade and at a new energy frontier. The LHC machine itself delivered above expectations, both during the proton-proton run that lasted from the end of March until early November, as well as the lead-lead period from mid-November until early December. At the International Conference on High Energy Physics (ICHEP) in Paris in July, senior physicists praised the grid for its ability to allow data to be analysed and presented in record time – just a matter of days as opposed to weeks or longer. This success is even more impressive if one takes into account the data volumes, rates and usage statistics that were all significantly higher than foreseen. With relatively few exceptions, the service was stable and available around the clock to researchers worldwide.

At the same time, the tools and services used by the Large Hadron Collider (LHC) experiments also benefited many others – from closely related activities, such as the Linear Collider Detector studies (featured as a case study in D6.1 “Capabilities offered by the HUCs to other communities” [D6.1]), to Life Science and Earth Science projects, which benefited not only from procedures developed and tested initially for the HEP community but also from specific tools and services.

Furthermore, the experience gained in processing and analysing the 15 petabytes (PB) of data acquired during this period (times a factor to take account for the replicas) have led to concrete “demonstrators” of optimised solutions for data placement (with greater emphasis on caching rather than static data placement), to shared tools for analysis as well as a roadmap for improvements in monitoring of the infrastructure and of the workflows. Coupled with a constantly changing environment – with virtualisation, multi-core, clouds and other technologies gaining increasing ground – there are many opportunities for further work in the next stages of the project.

The work supported by EGI-InSPIRE has been fundamental to these successes and has provided the necessary impetus to achieve common solutions in areas where these had previously been considered impossible. This clearly has consequences for long-term sustainability – both during the remaining lifetime of the project as well as beyond.

3.2 *Strengths, Weaknesses, Opportunities and Threats (S.W.O.T.)*

The following have been identified as the key strengths, weaknesses, opportunities and threats of the Services for HEP activity. Some are generic – such as the opportunity for increased commonality across disciplines – but are discussed below primarily in the context of this community.

- Strengths: 24x7 computing at the petascale;
- Weaknesses: operational costs – still too many problem incidents not resolved within 24h;
- Opportunities: increased commonality building on recent successes;
- Threats: failure to adapt to changing or emerging technologies.

3.2.1 Strengths: Production Computing at the Petascale

The LHC machine – two 27km rings that lie 100m underground – is cooled with superfluid helium to 1.9^0K and accelerates counter-circulating proton (or heavy ion) beams in an almost perfect vacuum to close to the speed of light before they are brought into collision in the centre of 4 massive detectors. After several levels of filtering and selection, the data that result from these collisions are processed and stored on the Worldwide LHC Computing Grid (WLCG) and analysed by scientists around the globe. The numbers involved are impressive: more than 10^5 CPU days per day, some 30PB of data annually (including replicas), global data transfer rates of several GB/s, in excess of 1M jobs per day and about two thousand unique analysis users per month (summed over the 4 main experiments). At the same time, the total number of GGUS tickets that are issued is typically around 100 per week – less than the total number of sites involved. The system delivers a high level of service 24 hours a day throughout the entire year (with somewhat lower service targets for 4 days around the end-of-year holidays). Not only has this service stood up well to the first year of data taking at the LHC but also it absorbed load that was higher than foreseen in many respects and critically in total data volumes and rates.

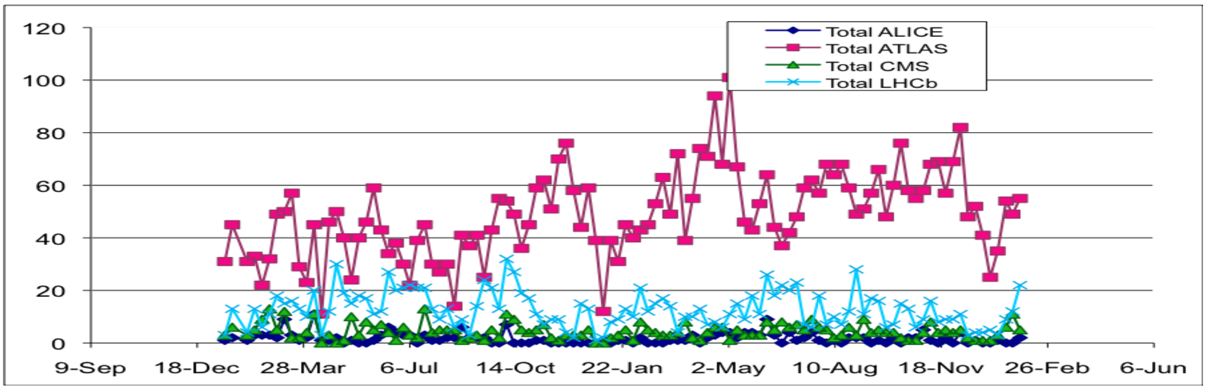

Figure 1 – GGUS tickets per LHC VO[1] [2]

3.2.2 Weaknesses: Operational Costs

Despite these impressive results that have been noted in the quarterly reports of EGI-InSPIRE, a number of service incidents take longer – sometimes significantly longer – than the targets to be resolved. Since the beginning of 2008, WLCG have required a Service Incident Report (SIR) to be produced for any major degradation or interruption to a WLCG service. An analysis of these reports highlights two areas where significant improvement could be realised at relatively low cost:

1. Network degradations (complete outages are typically followed up well whereas for degradations there has been some ambiguity regarding roles and responsibilities). A concrete example is documented in GGUS ticket 61440[3];
2. Database services: a significant number of prolonged downtimes could have been avoided had database recovery been swifter (and preferably performed either from an on-disk or else using a standby database). A concrete example of such an incident is documented in https://twiki.cern.ch/twiki/bin/view/DB/PostMortem18Dec10ATLARC.

[1] CMS uses GGUS only for reporting incidents to Tier 1 sites, hence the lower than expected numbers for this experiment.

[2] ALICE is shown here for completeness. Note that ALICE support was not EGI-funded during project year 1.

[3] *https://gus.fzk.de/ws/ticket_info.php?ticket=61440*.

In addition, storage services are still not sufficiently reliable, sometimes resulting in lost or corrupted files. The procedures to recover from such situations are typically very manpower intensive and often have to be tailored to the specific situation.

Why does this matter? No one has yet managed to improve on the wit and accuracy of Leslie Lamport's description of Distributed Computing[4]. Service degradation and/or outages can have an impact across the entire grid, potentially leaving thousands of scientists stranded – unable to analyse expensively acquired data. If service is resumed within one shift or even a day this can perhaps be accepted. If the interruption lasts longer – as is sometimes still the case – it surely cannot be tolerated and steps need to be taken to avoid such occurrences.

3.2.3 Opportunities: Further Improvements in Commonality

The first phase of the project has seen a number of notable achievements in the area of commonality, particularly in the areas of analysis tools and data management that are documented in more detail below. Recent workshops held by ATLAS and CMS have included sessions to identify further such areas and featured invited talks from the other experiment. Whilst any changes would need to be introduced in a manner that did not disrupt the on-going activities, there is significant scope for much larger changes in the future. However, further improvements in commonality are seen as highly desirable particularly as regards on-going maintenance and support costs.

3.2.4 Threats: (Not) Adapting to New Technologies

We are now in the third decade of LHC computing: the first (1990s) being characterised by the move to object-oriented languages and distributed computing, a series of research and development projects, the Computing Technical Proposals of the LHC experiments and the definition of the MONARC (Models of Networked Analysis at Regional Centres for LHC Experiments) [MONARC] model which defined the Tier0/Tier1/Tier2 hierarchy that is still in use today; the second (2000s) being that of grid adoption and wide-scale deployment; whereas the current decade (2010 on) will no doubt be accompanied by at least as many major changes in technology. As the computational and storage requirements of the LHC experiments will continue to grow, it will be impossible to ignore major changes in or evolution of technology. This, however, has a non-negligible cost, both in initial R&D as well as subsequent integration and production support, together with non-trivial migration costs. There is on-going work in a number of areas – adapting to virtualisation, investigating the use of clouds and/or cloud interfaces – that are relevant in this respect that need to be carefully followed.

3.3 Analysis Tools and Support

3.3.1 CMS

The CMS Remote Analysis Builder (CRAB) was the first analysis tool in CMS to aid users in configuring CMS applications for distributed use, by discovering the location of remote datasets and submitting jobs to the grid infrastructure. CRAB has progressed from a limited initial prototype nearly 5 years ago to a fully validated system that 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] estimated roughly 100K grid submissions per day. During the second half of 2010 the job submissions routinely exceeded the

[4] *"A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable."*

estimate by 40-50%. Despite this, CRAB coped well. After the development of the CRAB submission client, a CRAB server was developed, which has increased the scalability of submission and added capabilities of automatic resubmission. The CRAB server also provides a development platform for additional capabilities. While maintaining the current implementation the development team is in the process of reorganizing the architecture of the tool with the aim of reducing code duplication and increasing maintainability and sustainability.

During the past months four releases of the CRAB 2 Client and two for the CRABServer were produced. The main fixes have been applied to the data discovery and the splitting part of the workflow. All known problems with user data stage out have been solved. Problems with growing data set have been addressed and improvements in the management of the data masks have been applied. In order to improve the user support a new feature has been included. It integrates the Error Reporting Tool (ERT) designed to upload job state and logging information to a remote server. This has been done to allow analysis support teams to get detailed information about user errors allowing the effort needed by the end user support to be reduced. ERT is a component common to both ATLAS and CMS.

The main development tasks have been focused on the improvement of the interaction with grid middleware by reviewing the schema of the local database used by the new CRAB generation. It has been re-designed to keep track of job information mainly for logging purposes. A major reorganization of tables and field has been done in order to reduce the load at scale. Also the API structure, used to interact with the actual plugins that interface to the real scheduler, has been reviewed. In order to build a thin layer that is easy to maintain the plugin structure has been reduced to an essential set of calls.

Several tags of CRAB 3 have been prepared and related testing cycles have been done in collaboration with the CMS Integration team. The Credential API has been fully developed and integrated within the BossAir framework (the CRAB component that interfaces with the grid middleware). The role of the Credential API is to allow the support for multiple users. BossAir has been enabled to support python multiprocesses in order to optimize the performances. Plugins for the actual scheduler specific implementation (gLite, LSF, condor...) have been prepared. BossAir is now fully integrated and validated. A completely new strategy for the CMS user output data stage out has been also proposed and a work plan has been presented to CMS. The first prototype is now in progress as well as the corresponding monitoring tools.

The main items of development for 2011 will focus on the RESTful web-services in order to enable the CRAB3 Client-Server interactions. REST has been adopted because of its ease-of-use and to obtain full compatibility with the core CMS workload management services (which already communicate via RESTful web services). The aim is to deliver a prototype to be tested by CMS integration with a prototype for the client that must support the 4 basic functionalities: job creation, submission, checking of status and output retrieval. The development of the Asynchronous Stage out tool for user produced data handling is another high priority item. The plan is to couple this tool with CRAB3 server, but keeping it as an independent tool. A third goal is the definition and the implementation of the job sandboxes distribution strategy including the evaluation study of performances of both HTTPS and GridFTP. After collecting feedback from beta tester the aim will be at consolidation of the client and the RESTful interface, implementing additional functionality: kill, postMortem. Finally the developers want to implement the support for the automatic job resubmission. A pluggable agent is needed to offer the possibility of implementing use case specific algorithms for job resubmission.

3.3.2 Ganga/DIANE

Ganga is a frontend for job definition and management that provides a uniform interface across multiple distributed computing systems. Ganga is the main end-user distributed analysis tool in ATLAS and LHCb experiments, and it is also used as the foundation layer for other services such as HammerCloud to manage large numbers of jobs (see HammerCloud statistics). Ganga is a widely used tool outside of HEP[5]. From May 2010 to February 2011 Ganga was deployed in 127 sites and was used by 1316 users (40% ATLAS, 40% LHCb, 20% others) to submit more than 250,000 jobs weekly.

DIANE is a lightweight task processing framework which allows for more efficient and robust execution of large numbers of computational tasks in unreliable and heterogeneous computing infrastructures. Both Ganga and DIANE are used by power-users in various communities both within and outside HEP for running large-scale computing tasks and the use of these tools has been reported by communities from more than 10 scientific fields and disciplines. From May 2010 to February 2011 DIANE was deployed in 23 sites and was used to process more than 600,000 tasks.

Besides the LHC communities, the use of Ganga/DIANE tools has been reported for UNOSAT applications, Geant4 medical and space simulations and grid-enabled regression testing, AvianFlu Drug Search, ITU digital broadcasting planning, LatticeQCD simulations, Fusion, Image Processing and Classification, EnviroGrids, Simulation of Gaseous Detectors (GARFIELD).

3.3.2.1 Project Quarter 1

During PQ1 the work focused on software maintenance and evolution according to evolving requirements and features requested from the LHC user communities (ATLAS and LHCb). The work resulted in 10 public releases of Ganga, from version 5.5.5 to version 5.5.10-hotfix2. As stated in the Project Quarterly Report 1, during this quarter there was no funded effort to work in this area.

3.3.2.2 Project Quarter 2

During PQ2 the work focused on features to improve usability and user support. Notable new features include a web-based monitoring interface (WebGUI) to allow users to conveniently view the status of their submitted jobs and browse the local job repository, an integrated Error Reporting Tool (ERT) which uploads job state and descriptions to a remote server to enable user support teams to get detailed information about user errors, and finally Usage Monitoring has been improved to gather information about all submitted jobs regardless of submission backend – this is useful to understand the usage patterns outside of the known grid use-cases.

The Ganga WebGUI builds on a common web application framework developed for Task Monitoring Dashboard for ATLAS and CMS experiments. It allows the users to easily navigate between the job-repository view and central dashboard services, sharing the same look-and-feel and thus improving usability and user experience. The common web application framework allows support for monitoring applications to be streamlined, including LHC Dashboards, the lightweight Ganga/DIANE job monitoring service for other communities and the Ganga WebGUI.

The work also included evolution of experimental plugins for ATLAS and LHCb.

In this quarter the work resulted in 10 public releases of Ganga, from version 5.5.11 to version 5.5.16.

[5] See ganga usage monitoring at http://gangamon.cern.ch/django/usage

3.3.2.3 Project Quarter 3

During PQ3 the work in Ganga Core was focused on improving job merging and resubmission features. The framework now supports configurable auto-resubmission of failed sub-jobs and the possibility of overriding backend parameters when job resubmission is performed manually. The automatic merging code base has been fixed to ensure consistent location of merged outputs (which was not the case of Athena-based applications). The framework now supports job submission in the (optional) “keep going” mode, to enhance support for a large number of subjobs (hundreds or thousands). The support for job slices has been improved such that a slice may be constructed from an arbitrary list of jobs. Compatibility problems with newer version of python (2.6), batch backends (Sun Grid Engine) and grid middlware (gLite) have been fixed. GangaService package that is a generalisation of ATLAS skimming service implementation, has been added for possible reuse by other VOs. ATLAS Task Monitoring Dashboard plugins have been developed and put into production. Improvements have been made in the Ganga usage monitoring service, which now reports now the subjobs count and allows to better analyse the VO use of Ganga. A ramp-up in usage of the Ganga-based ERT use has been observed in CMS, with 200 reports uploaded in the period of 6 weeks since the CRAB 2_7_6 release. Restructuring of Ganga documentation is in progress and is nearly completed for the development wiki pages.

The work also included evolution of experimental plugins for ATLAS and LHCb.

In this quarter there were 8 public releases of Ganga, from version 5.5.17 to version 5.5.21.

DIANE has been used successfully in Q3 for the Geant4 regression testing with EGI and OSG grids. Some minor bug fixes and improvements to the mini-Dashboard task monitoring have been implemented and released. The project code repositories have been migrated to SVN. This work resulted in 2 public releases of DIANE, versions 2.2 and 2.3.

3.3.2.4 Current Work and Outlook for 2011

The development efforts will focus on making more efficient use of the EGI resources to improve task-processing throughput, and on tuning the user experience in order to decrease the user support load for the tools. Recent experience with Ganga-based ERTs for ATLAS, LHCb and CMS will be taken as a starting point to achieve better integration of the job submission tools with the user support systems in the future.

Another important work area is the improvement of the monitoring of analysis jobs (such as the ATLAS Task Monitoring Dashboard) so that grid faults are more quickly identified thereby leading to a lower failure rate of user jobs. In addition, support for intelligent fault detection may be required as an improvement for currently used distributed analysis tools.

Ongoing refactoring and integration of community-specific Ganga plugins (such as ATLAS tasks) into the common, core framework is another crucial point. Streamlining of features, such as the configured state for Ganga jobs, should help to reduce the maintenance effort and achieve more consistency of the interface exposed to the users.

3.3.3 HammerCloud

Figure 2: Concurrent HammerCloud jobs from May 1st 2010 to February 14th 2011

Since the beginning of the EGI-InSPIRE project, the HammerCloud (HC) distributed analysis testing service has been actively used to:

- continually validate the availability and performance of grid sites with short (~5-10 minutes), but realistic, jobs
- deliver on-demand stress tests using real jobs drawn from the user communities to the sites to aid in the commissioning of new sites or evaluate changes to site configurations.

Figure 2 summarizes the HC activity over the course of EGI-InSPIRE for the most active VO, the ATLAS experiment. ATLAS has run an average of 1000 concurrent validation and stress-testing jobs over the period; spikes in the plot above represent large multi-site stress tests.

The primary development goal for HammerCloud during this period has been to generalise the service so that it can be used by LHC experiments other than ATLAS, as well as VOs from other HUCs. This generalisation has been a success; HammerCloud plugins for the CMS and LHCb experiments have been developed and delivered during PQ1 and PQ2.

The recorded number of HC test jobs submitted over this period is summarized here:

- ATLAS # jobs (2010-05-01 to 2011-04-30) = 5280874
- CMS # jobs (2010-10-11 to 2011-04-30) = 3440394
- LHCb # jobs (2011-10-21 to 2011-04-30) = 14989

(The reported periods for CMS and LHCb are partial because records from the development/integration/testing periods and HC versions prior to version 4 were not kept.)

Details about developments during each quarter and current and future plans are given below.

3.3.3.1 Project Quarter 1

During PQ1 work was focused on the development of the CMS and LHCb plugins for HC. In fact this work began earlier in 2010, first with the development of a Ganga-CRAB adapter that allows CMS CRAB sessions to be managed via the Ganga API. By the end of PQ1, the HC-CMS code was completed and a trial deployment was made available for CMS users to test the service and provide feedback. The LHCb plugin development began in PQ1 and a prototype was rapidly made available thanks to the maturity of the existing Ganga-LHCb modules (Ganga is the official LHCb end user tool for submitting grid jobs).

3.3.3.2 Project Quarter 2

For CMS, PQ2 focused on the migration of HC service from a shared server intended for trial usage onto a dedicated submission host. After this migration, scale testing was performed in order to ensure the service could deliver meaningful stress tests to the sites. By working with the CRAB developers the HC-CMS plugin was improved to achieve higher job submission rates. In preparation for taking over site validation responsibilities from the CMS JobRobot, HC started running frequent short validation jobs on all of the CMS sites.

For LHCb, the HC prototype plugin underwent further development and by the end of PQ2 was ready to be deployed and tested at a larger scale with real users. While in the prototyping phase, the HC-LHCb service was employed to help validate a CASTOR upgrade at RAL.

3.3.3.3 Project Quarter 3

During PQ3, the CMS instance was stabilized while running the frequent validation jobs. The deployment of HC-CMS was considered complete by mid-PQ3 and a final presentation of the service was made to the CMS experiment. At the end of PQ3, CMS was in the process of identifying operations resources to complete the replacement of the CMS JobRobot with HC and to further integrate HC into their daily grid operations procedures.

For LHCb, the HC service was deployed onto a dedicated submission host, scale testing was performed, and the plugin development was considered complete. Further deployment of the HC-LHCb service and integration into the daily LHCb operations was postponed until PQ4. LHCb hired a former HammerCloud developer (to start in PQ4) to take over the deployment of the HC-LHCb service.

The developments in PQ3 culminated in the release of HC version 4. HCv4 introduces a system of "experiment applications" which are composed of modules to override the HammerCloud functionality during test submission, running, and presentation. This modular approach makes HC more amenable to new VO plugins, and the CMS and LHCb deployments were upgraded to HCv4.

The ATLAS HC instance remained at version 3 due to an increased user activity during the LHC winter shutdown; a version 4 instance was deployed, is under test, and will be put into production in PQ4. During PQ3 a prototype automatic-site exclusion extension was developed for ATLAS. This extension monitors the results of selected functional tests and sets the PanDA[6] queue status offline for malfunctioning sites. Lastly, development started on support for the testing of PanDA Production queues in order to validate Athena releases and PanDA pilot software for Monte Carlo production and data reprocessing activities.

3.3.3.4 Current Work and Outlook for 2011

Immediate HammerCloud developments for ATLAS are focused on three areas. First, the automatic site-exclusion extension is being completed, including an evaluation of various exclusion policies. Second, the support for submitting test jobs to PanDA production queues is being developed. HC is built around Ganga, which normally does not support job submission to production queues; Ganga development for production queue support has begun and an initial prototype was tested end-to-end with HC early in PQ4. Third, the ATLAS HC instance is being migrated to HCv4.

[6] Production and Distributed Analysis System developed by ATLAS, in part as an alternative to the gLite WMS.

For the remainder of 2011, it is expected that the CMS and LHCb instances will brought to the same service level as the ATLAS instance; adequate operational support from the experiment communities will be the key factor in the success of the HC integration with their daily grid operations.

Further planned core HC developments include an alarm system that detects and identifies specific problems by comparing performance metrics with their expected values, and in mid-2011 the database backend will be ported from MySQL to Oracle. High performance non-relational databases such as Cassandra may also be investigated.

3.4 Data Management Tools and Support

3.4.1 ATLAS Distributed Data Management

ATLAS fully relies on the usage of grid computing for offline data distribution, processing and analysis. The ATLAS Distributed Data Management (DDM) is a project built on top of the WLCG middleware that is responsible for the replication, access and bookkeeping of the multi-Petabyte ATLAS data (see Figure 3) across more than 100 grid sites. In this figure it is easy to locate events such as STEP'09 (large scale stress test of the WLCG infrastructure before data taking) or the start of the data taking activity (plus a clean up activity that occurred around December 2009). (Figure taken from the DDM Accounting pages).

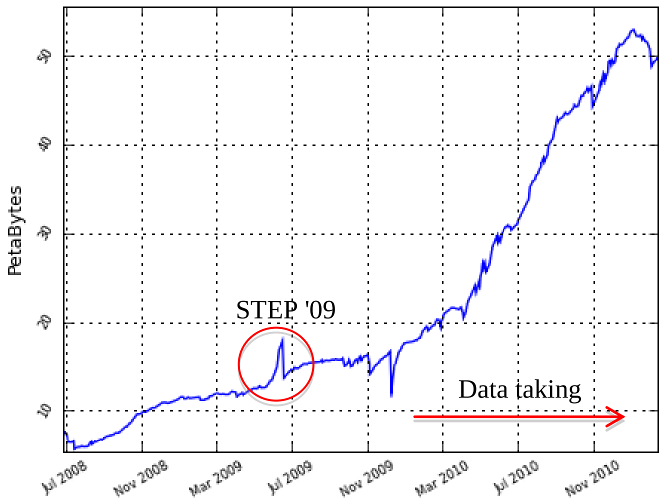

Figure 3: Evolution of the data volume managed by ATLAS DDM at all affiliated sites

In terms of support to ATLAS DDM, the work during the year was focused on:

- Facilitating daily operations work by improving the existing monitoring tools and implementing automatic tools to reduce the need of manual operations
- Optimizing the use of storage and network resources by adapting the necessary DDM components to the upcoming dynamic data placement models.

Apart from the usual background support to the operations team, the main development activities carried out each quarter are summarized below.

3.4.1.1 Project Quarter 1

3.4.1.1.1 Work on the DDM Accounting system

For over 2 years, ATLAS DDM has held an accounting database with historic space information for all ATLAS grid sites. This database is used to compare the storage usage reported directly by the storage elements (collected through SRM) with the usage known by the DDM Central Catalogues.

During the first quarter, a new agent was put in place to also retrieve the storage space information from the BDII. The information provided by the BDII is supposed to be more complete than that provided by SRM, as it should give for example an estimation of the space related to offline pools. From the first observations the publication of storage space to BDII seems to be less reliable and a simple web table was implemented to follow up and understand the differences between the SRM and BDII values.

The results of this work were presented in the ATLAS Software and Computing Week in July 2010 and were later taken to the WLCG Grid Deployment Board in September 2010. As a consequence, the use cases and needs of the different LHC experiments are being collected. They will appear in the "WLCG Information System Use Cases" document.

In addition to the above, new monitoring plots and views were developed to visualize associated metadata (i.e. data custodiality, needed for the evolution of the data placement model) extracted from the DDM Central Catalogues.

Given the continuous effort to improve the pages, the DDM Accounting has become a very popular monitoring tool in ATLAS distributed computing and is being used on a daily basis by the different site and cloud support teams as well as the central operations teams (see Figure 4).

Figure 4: Map overlay in the period from Jan 12 2011 to Feb 12 2011

3.4.1.1.2 Work on the ATLAS Centralized Blacklisting

The DDM Centralized Site Exclusion is a central database where the different DDM components can query which sites should temporarily be excluded because of downtimes or other conditions. Until Project Quarter 1, entries were made manually (by shifters and experts on call) which was a time consuming and error prone activity. In order to improve this situation two automatic collectors were deployed that are capable of setting sites offline. The first collector gets the official GOCDB and OSG Information Management System (OIM) downtime information from the ATLAS Grid Information System (AGIS) and is thus capable of excluding sites during the duration of their downtime. (AGIS aggregates the downtime information in a single service and provides an API to query this information.) The second collector allows ATLAS DDM to temporary stop replicating datasets to grid spacetoken endpoints with no free space.

The excluded site information is being published on a DDM webpage and fed to the Site Status Board in order to combine multiple sources of information and provide a single entry point for ATLAS users.

3.4.1.2 Project Quarters 2 and 3

The DDM Site Services are a critical service in ATLAS that is responsible for the data placement and throttling of the underlying gLite middleware. During the second and third quarters several new versions of this component were released and deployed. The main work was focused on reducing the ATLAS Cloud boundaries and evolution of the Computing Model. Following the ATLAS Computing Model, sites are grouped into ten clouds by geographical and organizational reasons. Each cloud is formed by one Tier1 that must provide a high level of service and several Tier2s and Tier3s, which depend directly on the Tier1. Given the network topology that was laid out for the data distribution, Tier1s usually act as the access point to get data in and out of the cloud and direct cross-cloud communication between Tier2s is generally avoided. However, since networking capabilities have significantly evolved, the Computing Model is driving towards more dynamic data placement policies where the replication of datasets is based on its usage rather than on static pre-placement. For this model to work optimally it is important that the cloud boundaries are gradually reduced.

Consequently, the DDM Site Services were instrumented to measure the durations of gLite File Transfer Service (FTS) transfers between sites and store them in an Oracle database. The transfer statistics are used as feedback to optimize the source and path selection particularly for cross-cloud transfers. The statistics are also visualized in a dynamic web page (http://bourricot.cern.ch/dq2/ftsmon/) in order to monitor the throughput performance of the network links.

In parallel, an ad-hoc load generator triggers transfers on the complete mesh of ATLAS sites. The results are summarized in the Sonar table (http://bourricot.cern.ch/dq2/ftsmon/sonar_view/cached/), which allows the different teams (central and regional operations teams) to filter the transfer statistics of the links that they are responsible to follow-up. The Sonar load-generator and its monitoring table are essential tools for the link commissioning activity of the ATLAS Grid.

The releases included other important features and corrections, which can be checked in more detail in the second Quarterly Report (EU Milestone: MS 106).

3.4.1.3 Current Work and Outlook for 2011

Apart from the general consolidation and support of the existing DDM framework, the plans for 2011 include the continuation of the work in the commissioning activity of the full mesh of site-to-site links and the tuning of the source and path selection mechanisms according to the acquired experience.

Another area of work is adapting DDM Site Services for Tier3 analysis sites. While these sites are allowed to have local control of their own resources, it is of general interest to have a unified solution to make the data available to them. Therefore DDM Site Services will be extended to consider particular requirements and generate new plug-in libraries that address the data discovery for these particular sites.

Finally it should also be mentioned that the improvement of the service level monitoring in DDM is already under way and should be ready by the next quarter: The monitoring infrastructure that is being instrumented will allow the early detection of system failures and service degradations.

3.4.2 DIRAC Data Management System

DIRAC system was developed in order to provide a complete solution for using the distributed computing resources of the LHCb experiment. DIRAC provides a complete framework for data production and analysis, including workload management, data management, monitoring and accounting. One of its most important components is the Data Management System (DMS), whose support from the EGI-InSPIRE project started in October 2010, therefore this report will include only the work done starting from PQ2.

3.4.2.1 Project Quarter 2

The activity of support for DIRAC DMS started in the last month of PQ2, and was initially focused in understanding the general structure of the framework, studying in detail all the services and agents which contribute to its overall functionality. The lack of overlap with the previous main developer (who left in October) and also of documentation made this task more complicated and time consuming than expected.

However, in parallel with the process of learning, a lot of effort was put in producing some documentation, in order to avoid such problems in the future.

3.4.2.2 Project Quarter 3

During PQ3, still a lot of effort was dedicated to the study of the system and its several components, especially the ones related to data management. In parallel, also significant progress has been done during these months in the improvement and development of the system. In more detail:

- Documentation for new developers has been produced. To do this, it was necessary to agree with the more experienced developers on a standard procedure to set a development environment, which had never been clearly defined before. This is the first guide for new DIRAC developers and will be certainly helpful for the next members who will join the DIRAC developer community.
- Enhancement of the functionality of the on-line database monitoring to visualize the status of the data while they are being transferred from the on-line storage system to the mass storage system at CERN. The service has been modified in order to display some more useful parameters, like the magnetic field state and beam luminosity. Integration with the web portal of the monitoring is ongoing.
- Development of a new DIRAC agent to allow consistency checks between the content of grid storage elements (SEs) and the information registered in the central file catalogue. Since the grid SEs and the file catalogue are completely decoupled, inconsistencies often arise, which have to be periodically fixed. So far this has been a manual, and thus error prone and time expensive, procedure. The new agent aims at making the full procedure in a totally automatic way. A prototype has been developed and is deployed in the development system. It is currently under testing and will move to certification during next quarter.

3.4.2.3 Current Work and Outlook for 2011

The main lines of development for this year have been defined as the following:

- Continue the development and integration of the new data consistency components to ensure the consistency of LHCb data stored in grid storage systems and the central file catalogues, started in PQ3.
- The LHCb Computing Model specifies keeping only the two latest versions of (re)processed real data on disk. However it is important to keep datasets that were used to produce published physics results for a very long time. A tape archiving mechanism, to be implemented later in 2011, is foreseen to address this use-case.
- Integration of new systems to provide dynamic data distribution in particular to monitor the usage of datasets over the grid storage systems and provide an efficient policy to remove less used datasets and create new replicas of the most used ones. This will bring to a more optimized usage of grid storage resources.
- Provide a system to keep account of the historical storage resources usage, grouping by different parameters (software version for the reprocessing, detector conditions, file type). This tool will display how much space a reprocessing or a new Monte Carlo production occupies and is urgently needed to manage storage resources at the grid sites in a more efficient way. This functionality is not provided by the basic middleware but is required by the VO to make efficient use of the storage provided. As described in the Description of Work, providing such missing functionality is within the mandate of this task.

3.5 Persistency Framework Tools and Support

3.5.1 Introduction

The Persistency Framework consists of three software packages (CORAL, POOL and COOL) which address the requirements of the heavy user communities in HEP for storing and accessing several different types of scientific data produced by the LHC experiments (ATLAS, CMS, LHCb). CORAL is an abstraction layer with an SQL-free API for accessing relational databases. POOL is a hybrid technology store for C++ objects, using a mixture of streaming and relational technologies. COOL handles the time variation and versioning of the conditions data of the HEP experiments. Two new personnel on EGI-InSPIRE funding, a PhD student and a fellow, joined the Persistency Framework team in June and July respectively. Their activities focused on R&D for data access optimization in CORAL and COOL for the former, and on the maintenance and consolidation of POOL and CORAL for the latter.

3.5.2 POOL maintenance and consolidation

In the area of POOL maintenance and consolidation, the main task consisted of the analysis and debugging of some problems observed during the automatic nightly builds and tests of the POOL software. Several bugs were identified and fixed and this task was successfully completed. First, software patches were prepared to fix some of the intermittent failures observed in the tests of the relational storage service component. As a second step, the configuration of POOL builds using the Configuration Manager Tool (CMT) was improved; this specific sub-task was also useful to gain a better understanding of CMT and of its usage in POOL. Finally, the conflicts resulting from the simultaneous access to the same database objects, by different tests running in parallel on different platforms, were addressed by implementing a locking mechanism to regulate the database access.

3.5.3 CORAL maintenance and consolidation

In the area of CORAL maintenance and consolidation, the main task was the analysis and improvement of the CORAL handling of network and database glitches. This is a high priority issue for all experiments, which have reported several problems of this kind accompanied by specific Oracle bugs. This task is still ongoing: while a workaround has been implemented and released to ATLAS and LHCb, a definitive solution to the problem still under development.

In a first iteration, this issue was addressed by its systematic study and resulted in a prototype patch to improve the stability by automatically recreating the physical connections and logical user sessions in the database.

In a second iteration during the next quarter, the code was reorganized to implement a strict check of the validity of the connection and session handles before each database interaction. After this code is fully validated and released, in the final phase a new strategy will be tried out. In this approach CORAL will better react to errors rather than only trying to prevent them.

3.5.4 R&D on data access optimization in CORAL and COOL

The R&D on data access optimization touched three different areas this year: the comparison of the FroNTier/Squid and CORAL server/proxy middle-tier and data caching technologies, an initial evaluation of commercial in-memory databases for data caching, and the usage of solid state disks (SSDs) for data access optimization.

The FroNTier/Squid and CORAL server/proxy technologies were studied by comparing their performance for CORAL data access retrieval to each other and to that of direct data retrieval from Oracle. Tests involved real use-cases (e.g. ATLAS High Level Trigger system) as well as simple data chunks of different types and sizes. The results are still being analysed.

The evaluation of the Oracle TimesTen in-memory database involved preliminary tests using simple custom data and native Oracle tools, which successfully showed a speedup of up to 5-6 times[7] in data retrieval when this was used as a caching layer. The integration of this technology with CORAL to retrieve the data of the LHC experiments, however, revealed issues that are not yet solved (missing support for some Oracle data types).

The more recent activity on SSDs focused on the preparation of a detailed plan to execute some tests when some dedicated hardware will be available in March or April 2011. With respect to hard disk drives, SSDs are faster, smaller, quieter, cooler and consume much less energy, but they are still much more expensive, so it is important to perform a cost-benefit analysis in the applicable areas. Three potential use cases have been identified: using SSDs to store the database tables containing conditions metadata; using SSDs as an additional layer for the internal database cache; using SSDs for undo logs or redo logs to speed up write access.

3.6 Monitoring and Dashboards

Monitoring of the distributed infrastructure and the activities of the user communities on this infrastructure is carried out in order to measure 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 VO activities, namely job processing and data transfer. The

[7] Preliminary PhD research results, yet to be published

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.

3.6.1 Experiment Dashboards

The system proves 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 5000 unique visitors (unique IP addresses) per month and more than 100K pages are accessed daily (the high number of accesses per user reflects the interactivity of the applications but may be inflated by auto-reloading views of the dashboard). 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.

3.6.1.1 Project Quarter 1

In the first quarter, a new job monitoring system was enabled for ATLAS jobs submitted via GANGA to the WMS. The original SAM system used by the LHC VOs was redesigned using the Nagios open-source framework for monitoring network hosts and services. In order to realign with the changes in the SAM architecture, a new Site Usability Portal started development.

3.6.1.2 Project Quarter 2

In PQ2 substantial progress was made in the development and deployment of the new version of the Dashboard's generic job monitoring application, which is shared by the ATLAS and CMS VOs. A new version of the historical view that shows job information as a function of time has been developed and initially deployed for the CMS experiment.

Most of the ATLAS jobs are processed by the PanDA workload management system. A new collector that enabled import of the job monitoring data from PanDA into the Dashboard schema was developed and put into production. prototype of a new version of the generic job monitoring for multiple execution back-ends and a new task monitoring application have been deployed to the ATLAS validation server.

3.6.1.3 Project Quarter 3

During PQ3 the main development effort was focused on job monitoring and Site Status Board applications. The Dashboard team supported the ATLAS VO in evaluating the Site Status Board application as a monitoring system for distributed computing shifts. The functionality of the Site Status Board application was extended and the user interface was improved.

The following improvements were performed in the job monitoring domain:

- The database queries for the generic job monitoring interactive view were optimised to improve performance.
- The generic job monitoring historical view was extended following the outcome of the ATLAS monitoring review to offer increased flexibility in the searching and filtering functions offered to users.

3.6.1.4 Current work and Outlook for 2011

During the last quarter the LHC user community evaluated new versions of the Dashboard job monitoring applications. ATLAS and CMS provided feedback and requested new functionality.

Following these requests the Dashboard team started to:

- Develop a new monitoring display for the production managers
- Modify the job monitoring historical application in order to add new distributions as for example number of processed data files as a function of time

The first prototype of the new ATLAS data management dashboard is being evaluated by pilot users.

A lot of improvements were performed on the Site Status Board application. Among them on the request of the CMS community a new algorithm for handling site downtime information was developed and deployed. The new algorithm takes into account topology information describing the sites and services used by a VO, retrieves information for scheduled and unscheduled downtimes from the OIM and GOCDB systems and properly handles changes of the scheduled downtime periods. The SSB collectors were partially redesigned in order to improve their performance and robustness.

At the end of 2010, the two biggest LHC VOs – ATLAS and CMS – held reviews of the monitoring infrastructure and tools used for their computing activities. The outcome of both reviews was positive for the Experiment Dashboard system and confirmed the important role of Dashboard applications in ATLAS and CMS computing operations. The outcomes of these reviews also defined the priorities for Dashboard development in 2011. During the next months of 2011 the Dashboard development will focus on the following directions:

- Job monitoring: the goal is to extend functionality and to improve performance, both of the data collectors and of the user interfaces.
- Data management monitoring: work on the new version of the ATLAS data management monitoring will continue and it should replace the current version by the end of 2011. The new version will allow users to follow transfer progress selecting either transfer source or destination and will facilitate detection of the transfer failures.
- Accomplish development of the new site usability interface compatible with the new SAM architecture.

3.6.2 Generic Dashboard

Though the target user community for the Experiment Dashboard system is the LHC, most of the Dashboard applications are generic and can be used outside the scope of LHC. Among those applications are Site Status Board, site usability interface and generic job monitoring; these applications are publicly available for other HUCs to install, customise and deploy. The applications like Site Status Board and site usability are mostly relevant for heavy user communities, that is those which use a lot of distributed resources and have special requirements for the distributed sites and services.

The data repository of the Dashboard job monitoring application is implemented in Oracle. This can create a limitation for using Dashboard job monitoring by smaller user communities. In order to satisfy the needs of smaller communities, the “mini-Dashboard” monitoring application was prototyped. It allows users to follow processing of user jobs submitted through the Ganga or Diane job submission frameworks and relies on MySQL for persistency implementation. The generic “mini-Dashboard” prototype is available at CERN:

http://gangamon.cern.ch/ganga

The “mini-Dashboard” installation at CERN provides basic functionality. The system can be extended and customized for the needs of a particular VO. It can grow together with the new user communities, by integrating their customizations and contributions if they are of general interest.

3.7 Summary of HEP Tools and Support

We have described above the progress achieved in supporting the HEP community in this critical year of the LHC start-up and of the transition from EGEE to EGI. Not only has this first year of data taking been very successful – quantitatively and qualitatively – but also there have been significant advances in commonality – including in areas that had previously been excluded (data management and analysis tools). A S.W.O.T. analysis of this activity has been presented, highlighting some areas where operational improvements can be made and the necessity of following the evolution of technology closely has been clearly identified.

4 SERVICES FOR LIFE SCIENCE

4.1 *Introduction*

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 Organizations (biomed, lsgrid, vlemed and pneumogrid). It receives support from 6 National Grid Initiatives (NGIs) (Dutch, French, German, Italian, Spanish and Swiss NGIs), the HealthGrid association[8], and two projects (LifeWatch ESFRI and EGI-InSPIRE). A board with representation for each stakeholder was nominated. To facilitate the user support activity of the LSGC, a set of user management tools is being developed with high priority during the first part of the EGI-InSPIRE project.

In addition, the Life Science community is involved in the development of the following community services:

- Dashboard: this activity is just starting in PQ4, at the time of writing this deliverable. The Life Sciences dashboard design integrates needs identified during the LSGC user management tooling development in order to integrate this information.
- Hydra encryption service: this activity is just starting in PQ4, at the time of writing this deliverable.
- GReIC: the GRelC RDBMs interface service has been under development throughout the reported period.
- Workfows: the work on Taverna workflow engine planned in this task to support the bioinformatics community has not started yet.
- CoreBio services: the work on bioinformatics databases provision has not started yet.

4.2 *Strengths, Weaknesses, Opportunities and Threats (S.W.O.T.)*

With EGI.eu, user communities enter a new era where they are supposed to self-organize and maintain their activity over time. The VRC model being set up in EGI-InSPIRE to liaise with users is still in its youth. Consequently, the SWOT analysis for Life Sciences activity is centred on the community set-up and its sustainability.

- **Strengths**: Ramping-up of the LSGC.

  The emergence of the LSGC and its growth since the start of the project is a very positive sign of the interest of the community for exploiting the grid infrastructure. The sustained activity of the LSGC-affiliated VOs is another measure of the stability of the LSGC community.
- **Weaknesses**: Multi-grid user community.

  The LSGC is a community of users exploiting diverse regional grid infrastructures, using different root authorities and middleware stacks. As a consequence, the problems faced by different members of the community differ significantly. Furthermore, interoperability between different infrastructures is hardly possible.
- **Opportunities**: Potential for community growth.

[8] HealthGrid association, http://www.healthgrid.org

There is a significant potential for LSGC growth. Several VOs gather communities from the Life Sciences domain, which are not affiliated to the LSGC yet. Some made the choice of creating independent VRCs but others are not affiliated to any VRC yet.

- **Threats**: No secure long-term funding.

  The LSGC depends to a large extent on the resources informally committed by several NGIs and organizations in terms of manpower. There is no guarantee of long-term resource commitment. Although the large number of contributor should ensure that no sudden halt of the LSGC is at risk, there may be significant reduction of the LSGC capability in the future.

4.3 LSGC user management tools

The LSGC user management tools include communications channels to animate the community, monitoring and user support tools, and tooling for managing the population of users registered in the community.

4.3.1 Communication channels

A mailing list of all LSGC partner representatives is used for internal communication. The LSGC has set up a wiki[9] to collect and publish practical and technical information related to the community. A monthly phone conference is organized to address the managerial and technical issues. The HealthGrid association maintains these communication channels that are essential in maximizing the impact of the tools and services developed, as well as negotiating and managing resources with the NGIs.

4.3.2 Monitoring and user support tools

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[10] 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 dedicated Nagios server[11].

4.3.3 Users management database

The LSGC is also currently designing a user management database, which will facilitate liaising with hundreds of users registered in the affiliated Virtual Organizations. The database schema was specified, and the tooling will be developed in the coming months. This user will interface to Virtual Organization Membership Service (VOMS) servers 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.

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

[10] Biomed VO, http://wiki.healthgrid.org/LSVRC:Biomed

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

Users will register to the “biomed” VO through a single portal collecting all information needed to fill-in the user database, register to the VOMS server and fill-in the application database. The user database will be interfaced to and monitored from the LSGC community dashboard.

4.3.4 Work accomplished

Project Quarter 1. The model for the LSGC was drafted during the first quarter of the project and endorsed at the HealthGrid 2010 conference held in Paris on June 28. This operational model defines governance policies and first actions including technical work on VO user administration tools.

Project Quarter 2. The LSGC wiki has been set up to collect and publish practical and technical information related to the community. A monthly phone conference starts being organized to address the managerial and technical issues. Technical work on VO administration tools started with the design of a VO users and application database and associated tools to monitor and manage the population of community members. The LSGC Technical Team of members from the biomed VO was set up. In QR2, the infrastructure monitoring is based on a lightweight Hudson integration server (http://hudson-ci.org/).

Project Quarter 3. The LSGC contributed to the identification and publication of stringent community requirements in collaboration with NA3’s User Community Board. The infrastructure monitoring was migrated from Hudson to Nagios. The user database design phase was completed with analysis of the “VO admin tool” capability and implementation of the file access control to the VO LFC (LCG File Catalog).

4.3.5 Current work and Outlook for 2011

Work is on going to replicate the critical biomed VOMS server and thus avoid that it becomes a single point of failure. The implementation of the user database just starts at QR4. A first prototype is expected in QR5 and a first complete version will be delivered in September 2011. The user database management tools will be interfaced to the LSGC Dashboard. The LSGC will also be active in supporting the Life Science community by negotiating Operation Level Agreements with EGI, the NGIs and the grid sites. This support will materialize in the establishment of Memoranda of Understanding.

4.4 GRelC data access interface

Support to the LS community in terms of grid-database management has been and will be provided through the Grid Relational Catalog (GRelC) software. General information about this service is available in Section 7.1.1

4.4.1 Work accomplished

Project Quarter 1. GRelC: Joint identification with the LS community of some domain-based use cases. New GRelC service instances are deployed to implement specific biological data sources.

Project Quarters 2 and 3. The SA3 questionnaire “A census about database resources, related needs and future plan” took place to assess the availability of databases that can benefit from the GRelC interface.

4.4.2 Current work and Outlook for 2011

The results to the SA3 questionnaire are being collected and analysed during PQ4. The results will be useful to better define and implement new use cases in the LS context.

4.5 Hydra encryption service

Work has just started during PQ4 on the Hydra encryption service for the life sciences community. The service was originally developed using gLite 1.5, so it will be evaluated if it can be deployed under gLite 3.2 without changes. After PQ4 a distributed and more secure service will be delivered. A functionality check of the Data Management System is then planned for every semester.

4.6 Issues and mitigation

The VOMS server and the LFC are VO-wide single point of failures that caused biomed VO-scale downtime due these services unavailability from time to time. A similar problem is likely to affect other VOs. Technical discussions are on-going on the best backup solution that can be proposed (e.g. using DNS aliasing to re-route request to a new server when a downtime is scheduled on the production one or duplicate the service and synchronize its database). The HealthGrid association will host a backup VOMS server.

Scheduled service downtimes are sometimes not properly reflected in the BDII, causing undue errors happening because the non-available resources are tentatively being used. Improvements to the usage of GOCDB downtimes in the BDII are expected in future.

4.7 Summary

During this first period, the Life Science community set up a Virtual Research Community, named LSGC. This community has defined governance rules and a line of action to facilitate access to the EGI infrastructure and improve experience of Life Science grid users. Technical services have been or are being developed to support this activity. During the next project years, more community services will be provisioned to enrich the portfolio of functionality available to build new Life Science applications.

5 SERVICES FOR ASTRONOMY AND ASTROPHYSICS

5.1 Introduction

Astronomers compose a worldwide, geographically distributed, community that operates large experiments and observing facilities. Astronomy, in fact, is an observations-driven science based on national and international ground- and space-based observatories producing terabytes of data, which are publicly available and stored in astronomical archives all around the world and usually made publicly available to interested scientists. Observations done at different wavelengths, from radio and infrared, through visible light, ultraviolet, and x and γ rays, are providing a new, panchromatic picture of our Universe, and lead to a better understanding of the observed objects.

The A&A community is involved in some innovative and technologically advanced experiments that are setting up new observational facilities: the A&A ESFRI related projects such as the future European Extremely Large Telescope (E-ELT), the Square Kilometre Array (SKA) and the Cerenkov Telescope Array (CTA) represent concrete examples of such facilities.

Considering the huge amount of data coming from past and modern observing facilities as well as those generated by simulations and those available in the Virtual Observatory and their computational needs, the use of DCIs with particular reference to those based on grid technology is of strategic importance for the advance of astronomical research.

5.2 Strengths, Weaknesses, Opportunities and Threats (S.W.O.T.)

The failure of the so far proposed projects to create a well organized A&A user community able to fully exploit DCIs (with particular reference to the EGI Grid DCI) has forced to scale down the community efforts. The A&A community has the potential to effectively exploit DCIs and their related tools and services; to achieve this objective, however, a big effort is requested to improve the coordination of the community and create a well organized VRC so that isolated efforts and initiatives can be profitably conveyed to provide value-added services for the EGI Grid-based DCIs and perform effective training and dissemination activities.

In the next paragraphs the SWOT related to the A&A community is reported:

- **Strengths**: The community has gathered a number of successful experiences with DCI e-Infrastructures in the past years; the participation in the EGEE projects and now in EGI-InSPIRE, as well as in the European Virtual Observatory, are among the most important examples of such experiences.

  The Virtual Observatory is composed of a set of federated, geographically distributed world-wide, major digital sky archives, equipped with the software tools and infrastructure necessary to combine them in an efficient and user-friendly manner, and to explore the resulting data sets whose size and complexity are beyond the capabilities offered by traditional approaches. The Virtual Observatory, therefore, is an excellent example of worldwide data e-Infrastructure in the Astronomical domain. Such data e-Infrastructure needs now to be complemented with DCIs to provide the necessary computing facilities.

  The A&A ESFRI related projects as well, namely the future European Extremely Large Telescope (E-ELT), the Square Kilometre Array (SKA) and the Cerenkov Telescope Array (CTA) greatly contribute to strengthen the community because: a) they are potential heavy users of DCIs and of their facilities; b) they could play an important role as aggregators of the whole community.

- **Weaknesses**: Astronomers form a complex community with several research branches active within it. Until now the attempts to coordinate the European A&A community have not produced the expected results. Many times different projects/institutions/research groups tend to set up and adopt their own tools and services whose coordination/harmonization is quite hard to achieve. The transition from EGEE projects to EGI and the concurrent rejection of the proposal aimed at creating the A&A VRC worsen this situation. To achieve a good level of cohesion in the community, it would be necessary to form a new project able to catalyse resources, e.g. a new dedicated project or some of the A&A ESFRI projects. In conclusion the A&A community (at least for what concerns DCIs and HPC) suffers from an excessive tendency to fragment and this is probably its major weakness factor.
- **Opportunities**: A&A community has the potential to intensively and effectively use DCIs. It is therefore a community that might grow more and more in terms of users (including ESFRI projects) and related applications that could take a great benefit from massively using DCIs.
- **Threats**: If no funds are made available in the near future to set up a VRC able to: a) provide those tools and services (including portals and science gateways) requested by A&A applications and projects; b) train its potential users so that they are enable to approach and use DCIs (especially the EGI DCI), it will be extremely difficult to overcome the excessive dispersion that currently plagues the A&A community. With no services and tools and an inadequate training/dissemination activity really able to overcome the barriers currently interposed between end users and DCIs the grid has no chances to be largely adopted by astronomers.

5.3 Work accomplished

5.3.1 Project Quarters 1 and 2

In these two quarters the A&A activity focused on management tasks and on preparatory studies. Topics identified include: Visualization Tools, Parallel processing on the grid and CUDA, access to databases and integration with the Virtual Observatory.

An internal work-plan was prepared to identify in detail the sub-tasks and the involved resources. Some preliminary activities on the A&A subtasks have been done. In particular:

- An evaluation plan was defined to verify the state of the art of Database support and possible integration with Astronomical archives and catalogues (stored in databases) and with the Virtual Observatory.
- A preliminary study was performed on VisIVO as well as a preliminary integration plan to verify grid tools and services already available that could be used to visualize data distributed on the grid.
- Collaboration in the framework of the MPI working group has started to collect A&A community requirements. The option to use CUDA for some A&A applications was evaluated as well as the current support for CUDA on the grid.

5.3.2 Project Quarter 3

The first fundamental step of the gridification of VisIVO aims at having movies and images directly stored on the grid, even if intermediate files are not produced, and at reducing the overall time for movie production.

For this reason two main steps are requested to complete the porting of VisIVO on a gLite-based grid. The first step is to allow the usage of VisIVO directly within a code during the production phase. We are now developing a software layer to use VisIVO directly within codes using the internal arrays without the need of producing intermediate files. A library of VisIVO was designed from November to December 2010 and now the development phase is in progress.

The second step aims at making possible the execution 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 cannot start before the deployment of the VisIVO Library in order to avoid any problem during the design phase of the new MPI-enabled VisIVO for the grid.

For what concerns grid computing and HPC, the second half of 2010 was spent identifying significant A&A use-cases and test-beds and planning a coordinated activity in the context of EGI.eu and of NGIs. Cosmological simulations represent one of the most important classes of A&A applications requiring HPC resources. We identified FLY (INAF-OACT Cosmological code) and Gadget + Flash, the most common cosmological codes in Astrophysics.

We are now in the process of collecting requirements from these applications following this schema: a) preparation of the initial dataset; its size is of several hundreds of Gigabytes; b) data production phase, generally performed through parallel code whose execution involves hundreds of CPU/cores. We are now starting the design of some preliminary tests to run in gLite.

The A&A community is now evaluating tools and services currently in place to integrate/interface grid infrastructures and databases to use them in the context of A&A applications. Tools and services currently under evaluation include AMGA, GRelC, Spitfire and OGSA-DAI.

A report will be shortly issued concerning all evaluated tools and services and the outcome of this evaluation process (those selected to be used for A&A applications). The report will also clarify the selection criteria.

5.4 Current work and Outlook for 2011

For what concerns VisIVO, the following steps will be accomplished during 2011: a) the VisIVO Libraries API (Application Programming Interface) will be completed; some use cases will be prepared and tested on the gLite infrastructure; b) prototypes for importers and filters will be implemented and tested; c) the MPI-compliant version of VisIVO will be deployed.

VisIVO Libraries represent an important step in order to export the usage of the 3D visualization rendering and movie creation tools to other user communities. When this phase is successfully completed, the MPI-enabled version of VisIVO will be tested on the infrastructure, thereafter it will be deployed.

The activity foreseen during 2011 related to grid and supercomputing includes 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). To this end we aim to involve a significant number of A&A research groups (also those not directly involved in EGI-InSPIRE activities). This objective would be greatly eased if the necessary additional resources are made available through the new funding channels identified during the last year. The accomplishment of activities related to grid and supercomputing requires a strong coordination within EGI-InSPIRE given that this is one of the key objectives of the project at large and transversal to many of its communities.

Integration between grids and databases remains 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 a chance of penetrating the A&A community. The evaluation of tools and services such as AMGA, GRelC and possibly others will continue.

New activities might be jointly carried out 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.

The identification of use-cases and test-beds well representing the majority of A&A applications is a key activity that, in turn, requires an adequate number of A&A institutes and research groups to be involved.

6 SERVICES FOR EARTH SCIENCE

6.1 Introduction

Earth Science (ES) applications cover various disciplines like seismology, atmospheric modelling, meteorological forecasting, flood forecasting and many others. The presence in SA3 is centred in the implementation and maintenance interfaces or tools to provide access to Earth Science specific resources from the grid, for example resources within the Ground European Network for Earth Science Interoperations - Digital Repositories (GENESI-DR).

The ES Grid community also includes researchers and scientists working in the climate change domain. The Climate-G testbed provides an experimental large-scale data environment for climate change addressing challenging data and metadata management issues. It is based on the Grid Relational Catalogue (GRelC, see Section 7.1).

6.2 Strengths, Weaknesses, Opportunities and Threats (S.W.O.T.)

The failure of the DCI user community related proposal for Earth Science (ES) has forced a reorientation and downscaling of the community efforts. The community is still alive and supported independently by organisations and NGIs, but effort is required to foster the community and to provide value-added services around EGI. The Services for Earth Science task only covers the implementation of data access scenarios, to permit the utilization of Earth Science data resources in grid jobs.

The following have been identified as the key strengths, weaknesses, opportunities and threats of the Services for ES activity.

- **Strengths**: Existing user community initiative.

  The community that was built around Earth Science centred research utilizing grid resources during the EGEE projects is still collaborating actively. The resources that are a result of this work can still be built upon now, without effective funding. This includes the support of the catch-all Earth Science Virtual Organisation for researchers not yet organized in specific VOs, which is still maintained and supported. The web portal http://www.euearthsciencegrid.org/ with its included collaboration tools can be used to communicate and organise effort.
- **Weaknesses**: Diversity of demands and technologies

  Earth Science research comprises many different disciplines with a wide range of technical requirements. As data is a key point for most of these disciplines, many different mostly separate systems, standards, projects and initiatives exist around data and metadata management. The Services for Earth Science task can only consider a fraction of these possible promising targets, and as such can only satisfy parts of the community.
- **Opportunities**: Potential for community growth.

  There is a significant potential for growing the community through simplifying the access of Earth Science data systems from EGI.
- **Threats**: Absence of funding

  The ES HUC highly depends on the resources committed by NGIs. The main threat concerns the organisation and community activity in medium term without any funding for networking. ES Grid teams exist in many European countries and are linked to EGI via their NGI. They have previously expressed the need to work together in a common framework in order to take full advantage of the scattered expertise but there is no guarantee of long-term commitment if none of the FP7 proposals submitted was funded.

6.3 Work accomplished

6.3.1 Project Quarter 1

A memorandum of understanding about the collaboration with the GENESI-DEC project, the successor / evolution of the GENESI-DR project for an open Earth Science repository infrastructure, was discussed. One of GENESI-DEC's major goals is to enhance the previously established platform by federating and interoperating existing infrastructures of Digital Earth and Earth Science initiatives. The major goal of the collaboration is to allow and ease access to the data infrastructure for users. Additionally, integration with existing tools commonly used in the grid by Earth Scientists was investigated.

GRelC has been used extensively in the Climate-G testbed, which is a multidisciplinary collaboration involving both climate and computer scientists and it involves several partners such as: Centro Euro-Mediterraneo per i CambiamentiClimatici (CMCC), Institut Pierre-Simon Laplace (IPSL), Fraunhofer Institut fur Algorithmen und WissenschaftlichesRechnen (SCAI), National Center for Atmospheric Research (NCAR), University of Reading, University of Salento and SPACI.

A new version of the GRelC service was released and it was tested to manage climate metadata. A new version of the Climate-G portal was deployed (http://grelc.unile.it:8080/ ClimateG-DDC-v2.0/) and tested.

6.3.2 Project Quarter 2

To ease the management of the software provided by SPACI in the context of the Climate-G testbed, a Virtual-Machine (VM) based environment has been set up. A new VM for the GRelC service (to manage the metadata information) is now up and running as well as a new VM for the Climate-G portal. The dashboard page of the Climate-G Portal has been improved. The update includes the “compare variable (advanced)” option. Summary data is now also provided.

New requirements and scenarios will be gathered and defined together with the Earth Sciences representatives involved into the EGI-InSPIRE project during PQ3. Finally, the VOMS service for the Climate-G VO has been put into operation and was tested successfully.

Possible solutions for the community to use an implementation of OGC Web Processing Service (WPS) for submitting jobs on EGI have been evaluated. The most important aspect of this work is the authorization and authentication. The WPS standard itself does not include authentication or authorisation capabilities, but different solutions have been developed in the past, e.g. for Globus from the GDI-Grid project (http://www.gdi-grid.de/index.php?id=69&L=1) of the German D-Grid initiative, or for gLite from the Open Grid Forum G-OWS Working Group (https://www.g-ows.org/). The community faces the same problem when trying to access OWS-compliant data sources, such as instances of WCS (coverage) or WFS (features). GENESI-DEC is also concerned with on demand processing capabilities, which will be interfaced by a WPS compliant service as well. A "federated security interoperability framework" such as described by GENESI-DEC, which would map the different security models of the selected infrastructures (GENESI-DEC inititially envisions EMSO, SeaDataNet/GeoSeas, CEODE and A-SDI) to allow secure access, would thus also be highly beneficial for the EGI user community.

6.3.3 Project Quarter 3

A closer collaboration with the ES Community in the French NGI was established. Together with the Institute Pierre Simon Laplace (IPSL) in Paris the interfacing of the EGI infrastructure with the Earth System Grid (ESG) was analysed and possible solutions were evaluated. It is planned to access a CMIP5 Data Node from EGI. The underlying technology here is based on OPenNDAP (see MS602 3.4). This project will be done in collaboration with other climate projects in France and in Europe. The work will be partly supported by non-EGI-funded manpower.

The use of the OGC (Open Geospatial consortium) components with gLite was developed during the CYCLOPS project. Two applications, one on fire modelling in Italy and Portugal, the other on flash flood modelling in France, have used this new capability. As geospatial components are frequently used in ES a tutorial was organized in France. A server was deployed and the application started gathering the data.

The “SA3 - Questionnaire – A census about database resources, related needs and future plan” was sent to the ES community. Preliminary feedback from this community is expected to be provided during PQ4. Moreover, in the context of the Climate-G test case, the Climate-G Portal was updated. Some bugs connected with the visualization of the climate change impacts on economy (histogram charts) were fixed.

6.4 Current work and Outlook for 2011

The Services for ES activity is centred on the access to ES Data infrastructures. However, the ES HUC of this WP also performs unfunded operational support for ES and identifies new topics that are of increasing interest for ES EGI users. This year, the community will work in the following areas:

- Work on the data access tools / interfaces in funded part of task Services for ES.
  - Analyse new features that are under development in the on-going GENESI-DEC project
  - Foster direct collaboration with GENESI-DEC
  - In collaboration with IPSL find a solution to access the Earth System Grid (ESG) data from EGI. Potentially implement a demonstrator use case as a proof of concept to stimulate the interest of the climate user community.
- Operation and maintenance of the catch all ES VO (ESR) will continue unfunded.
- Exploration of alternative sources of funding, to sustain the core of the VRC activities, to continue collaboration among the teams

7 OTHER COMMUNITIES & SHARED SERVICES AND TOOLS

7.1 Grid Relational Catalog (GRelC)

7.1.1 Introduction about GRelC and its role in EGI-InSPIRE

The management of databases plays a crucial role in different scientific domains: Earth Science, High Energy Physics and Bioinformatics are just some notable examples. In the last decade a lot of efforts have been devoted to the grid data management research topics, trying to design, implement or improve existing data services to manage heterogeneous, geographically spread and dynamic data sources in a grid environment. Started in 2001, the GRelC Project provides a grid-database service allowing users to interact with different database management systems, both relational (PostgreSQL, MySQL, Oracle, DB2, SQLite, etc) and non-relational (eXist, XIndice, XML flat files). It provides a uniform access interface to heterogeneous data sources in a grid environment. Since 2008, the GRelC middleware has been included in the EGEE RESPECT Programme since it works well with the gLite software by expanding the functionalities of the grid infrastructure (with regards to the grid database management).

During the EGI-InSPIRE project, the GRelC software will support the HUC communities with a new set of functionalities available as web application through the GRelC Portal and the DashboardDB interface. The DashboardDB will represent the proper web access interface to the EGI Database of Databases. Such a “registry” will complement the functionalities provided by the EGI Application Database and will represent a distributed and multi-VO system supporting the HUC. It will be searchable and will allow people 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. All of these Web2.0 functionalities will help in creating *communities* around the available grid-databases. Providing social-oriented functionalities along with the classical ones related to the access and management of the data sources is a key point for the next generation of grid-database systems.

In terms of support, the GRelC project provides a website (www.grelc.unile.it) where users can find information about the status of the project, new releases, installation guides, software development kit, rpms, etc. Support is also provided through some tutorials available on the GILDA[12] website where people can learn more about the GRelC service, the GRelC Portal, the Command Line Interface, and so on. Tutorials can be found on the GILDA Wiki and are organized to provide different skills according to different users’ needs and requests. Basically, the GRelC service provides support to the HUC in two different ways: in terms of *grid database management* and in terms of *grid metadata management*. For instance, in the Earth Sciences and Environmental context the main contribution concerning GRelC is related to the Climate-G testbed. The Climate-G testbed provides an experimental large-scale data environment for climate change addressing challenging data and metadata management issues. The main scope of Climate-G is to allow scientists to carry out geographical and cross-institutional climate data discovery, access, visualization and sharing. The grid metadata service of the testbed is the GRelC service and it manages both relational and XML metadata compliant with INSPIRE directives as well as with ISO standards (ISO19115 and ISO19139).

[12] https://grid.ct.infn.it/twiki/bin/view/GILDA/GRelCProject.

7.1.2 Strengths, Weaknesses, Opportunities and Threats Analysis (S.W.O.T.)

7.1.2.1 Strengths

- A Community-based approach can give the registry application (EGI database of databases) a great opportunity to attract new users and create ***communities*** around grid database resources. It can help users in knowing what other people are doing in the same/other scientific contexts, post comments, give suggestions, share results, etc.
- A ***login/password*** based authentication (like the one adopted in the most common social applications) will be able to attract a large number of users (even including non-grid users). Conversely, a certificate-based authentication is not high-priority or high-impact as the login/password one for this kind of (web) application.
- Exploiting ***permalink*** capabilities (to export specific web pages like the registry) can strongly address reusability in several web applications. This way the registry could be integrated into other web contexts in a straightforward manner.
- The registry will allow users to:
    - ***discover*** existing grid-databases resources;
    - ***publish*** new grid-database resources;
    - ***ask to get access to*** a grid-database;
    - ***join a discussion group*** related to a grid-database resource.
- Keeping so many important information, the registry can represent a notable ***knowledge base***.
- The DashboardDB will provide support also in terms of ***grid-database services monitoring*** making these services at the same time more visible to the whole EGI-InSPIRE community.
- Having a registry of databases can motivate people to:
    - keep ***up to date resouces*** that could be of interest for several users/communities;
    - provide ***high level of availability*** with regard to the grid-database scenarios.

7.1.2.2 Weaknesses

- There are no relevant weaknesses concerning the DashboardDB development. The registry implementation is going according to the initial roadmap and the first version of the registry and the DashboardDB container will be available in Q4 and presented at the EGI User Forum in April.
- No answers have been collected so far regarding the SA3 Questionnaire on grid database resources. The Questionnaire has been sent just at the end of PQ3 (January 16, 2011).

7.1.2.3 Opportunities

- The large and multifaceted EGI community can take advantage and be the primary user of the DashboardDB and the database registry. Potential users exist (i.e. bioinformatics, astrophysics, geoscientists) and the registry can represent a good place where ***promoting*** grid-database resources.
- The scientific community needs to know what people are doing in several countries or contexts, to start collaboration, ***find people***, exchange comments and opinions, ***share knowledge***, ***avoid the replication of work***, etc. A registry of grid-databases (as the one proposed in this activity) can support these tasks. The search & discovery functionalities jointly with the Web 2.0 approach address these challenging opportunities.
- There are ***well known database resources*** available through the Internet that could be ported in grid by the related user community and published on the registry. Grid users could benefit from the existence of these grid-enabled resources and new use cases could be implemented.

- ***Several tools and services*** (with different features and functionalities) help users in managing/porting databases in a grid environment and are already exploited by different HUC. This variety of tools is an opportunity to have more entries in the grid database registry.

7.1.2.4 Threats

- Several users exploit database resources in grid without thinking about publishing or promoting them in a more general context like the registry for further exploitations. A strong dissemination is needed among the HUCs to better understand the registry potentialities and benefits.
- A public registry is not only intended for public access databases. This is an important concept that the HUC has to be aware of. Grid-databases with strong policies and access restrictions can be published too. The registy does not grant any access to the target resources.

7.1.3 Work accomplished

7.1.3.1 Project Quarter 1

In the first quarter one person was hired to start working on the tasks related to the GRelC service. In particular a *system database* to support the management and monitoring framework of the GRelC services network was designed and implemented.

The main subtasks and achievements were the following:

1) identification of the most important classes of information related to management and monitoring activities. First version of the Entity/Relationship (E/R) diagram modelling such concepts and of the associated logical schema. First implementation (MySQL based) of the system database. Preliminary tests and bug fixing.
2) Extension of the system database to include a registry section. New classes of information concerning the grid-databases managed by the GRelC services were identified and modelled. An extended version of the E/R model, a new logical schema and the related MySQL implementation were provided. This work represents a core part of the EGI Database of Databases described in the Description of Work.
3) Design of the management and monitoring client performing the data ingestion into the system database. The client takes care of the databases and the GRelC services availability and status. It must check and store all of the information identified in the previous tasks. It is important to note that this client will be GSI and VOMS based to retrieve the relevant information about the available grid-databases, the supported Virtual Organizations, etc. It will be built on top of the GRelC libraries.
4) A new version of the GRelC service was released and tested to manage climate metadata.

In the Life Sciences context, the provided support has been related to a new GRelC service instance to implement specific use cases exploiting LS (e.g. biological) data sources.

7.1.3.2 Project Quarter 2

During Q2 three main tasks were performed:

1) a management and monitoring client (named DashboardDB Client) was developed. It acts as an information provider able to retrieve and store important metrics related to the GRelC services network. Available statistics are: RTT, availability/service down, network errors, host unreachable. The client can be configured through a configuration file to define the frequency of the measurements. Information about hosts and services are directly retrieved

from the system database developed in Q1. Several tests were carried out to evaluate the robustness and performance of this application. In the future (v2.0), the same client will be able to retrieve information about the grid databases managed by each GRelC service instance.

2) Preliminary design of the management and monitoring web interface. The design takes into account the hierarchical structure *project<->host<->service* discussed in PQ1. During this phase a complete analysis about the functional and non-functional requirements was carried out. The following non-functional requirements were considered crucial for the proposed web application: portability, look & feel, interoperability, and transparency. On the other hand the most important functional requirements that were identified during the design phase were: (i) project view, (ii) host view, (iii) service view, (iv) GRelC registry. The service view must provide a complete dashboard with charts and tables displaying: RTT measures, service availability/service down, summary information, local registry. Preliminary design includes also the class diagrams related to the “charts” classes. About 20 packages were identified during this phase to manage charts, beans, security, etc. The design will continue in Q3. Moreover, several software providing charts capabilities was tested to identify the ones suitable for this activity.
3) The training environment was tested and updated from several points of view: (i) Wiki, (ii) machines (one in Lecce and another one in Catania), (iii) available database resources, (iv) GRelC portal, (v) Command Line Interface.

Concerning the support to the Earth Science community, a Virtual-Machine based environment was set up. In particular:

- a new VM for the GRelC service (to manage the metadata information) was deployed;
- a new VM hosting the Climate-G portal was deployed.

The Climate-G Portal was improved in the dashboard page, in particular referring to the “compare variable (advanced)” option. Summary data are now also provided.

7.1.3.3 Project Quarter 3

The design of the DashboardDB (started at the end of Q2) and the implementation of some internal modules (Java classes) were the main activities carried out in the QR3. It is worth mentioning that the design of the internal modules of the DashboardDB web application performed during Q3 has taken into account the Web2.0 paradigm. Mashup, Google Maps, permalinks, comments, are just some of the features that have been considered during the design phase.

The DashboardDB design (PQ3 activity) implements the Model-View-Controller (MVC) design pattern. This way a clear separation of concerns allowing to manage the complexity of the web application is strongly provided.

A monitoring view is also part of the DashboardDB application and will give the users the proper understanding about the underlying grid-database service infrastructure. Such a view will be useful to monitor and check the network of GRelC services deployed within the EGI context. The main modules that were designed and developed during PQ3 are related both to the monitoring part and to the grid-database registry one. In particular, some packages that were designed and implemented during this period include:

- core package (abstract classes, common modules, utils routines)
- charts package (pie, bar, etc.)
- beans package (for business objects like GRelC services or grid-databases);

- stream package (Java classes producing data streams to export data in CSV format);

For each of these modules, several unit tests have been carried out (in Q3) to remove bugs and make the code more robust and stable.

Another important task carried out and completed during PQ3 was the "*SA3 - Questionnaire – A census about database resources, related needs and future plan*".

The questionnaire aims at providing an up to date list about databases (relational, XML-based, etc.):

- already in place but that need to be ported in grid in the context of the EGI-InSPIRE project;
- already ported in grid and so accessible in the context of the EGI-InSPIRE project;
- not yet deployed in the context of the EGI-InSPIRE project;
- available from external sources via FTP, HTTP and that would need a grid-enabled instance and interface in the context of the EGI-InPIRE project.

The questionnaire has been sent to the HUC at the end of the PQ3 (January 16, 2011) and preliminary feedback is expected to be collected during PQ4. Support to HUC (e.g. Life Science and Earth Science) will start from the output of the questionnaire and will be use-case driven.

Finally an abstract about the GRelC activity has been submitted and accepted for oral presentation at the EGI User Forum in April 2011.

7.1.4 Current work and Outlook for 2011

The implementation of the DashboardDB application will continue in PQ4. A preliminary version (v0.1) with the home page will be available at the end of PQ4 just to highlight the main goal of the DashboardDB and start collecting feedback. Some internal modules, such as the project and registry management, will be also developed during PQ4. A preliminary static list about the available database resources in the context of EGI will be inferred in PQ4, as a follow up of the SA3 Questionnaire sent to the HUC during PQ3. The list will help to define use cases, understand needs and collect requirements from the HUC.

During PQ5 the project and registry views will be completed. The project view will allow groups of GRelC services to be managed as a single logical entity. For instance there could be an "EGI" project including all of the GRelC services running in the EGI production grid. The registry (the EGI Database of databases) view will allow people to find/discover database resources, apply filters by tag, VO, publication date, etc. Moreover users will be able to join or leave discussion groups about these data sources, post comments, rate the grid-databases and add tags according to a community-oriented vision and a Web2.0 approach.

During PQ6 and PQ7 the host and service views will be almost completed. The host view will provide information about the host resources belonging to a specific project, whereas the service one will provide a service-oriented dashboard with charts and reports about service statistics and availability.

7.2 SOMA2

SOMA2 is a versatile modelling environment for computational drug discovery and molecular modelling (http://www.csc.fi/soma). SOMA2 is operated through a WWW-browser and it offers an easy access to third-party scientific applications. The SOMA2 environment offers a full scale modelling environment from inputting molecular data to visualization and analysis of the results, and including a possibility to combine different applications into automatically processed application workflows.

7.2.1 Work accomplished

7.2.1.1 Project Quarter 1

First milestone in SOMA2 development for EGI-InSPIRE was to extend existing SOMA2 gateway to interface with job submission to grid infrastructure. During the first project quarter we did design and targeted development to achieve the milestone. As a result, a support for processing user's personal X509 certificate was added in SOMA2. System is now also capable of examining existing certificate and if no certificate is found for a user, one can be attached into the system. Also, we are now able to define grid entities in SOMA2 program description scheme. Grid entities are integrated into existing SOMA2 service authorization scheme so that available resources (scientific applications) for a user can be resolved according to user's certificate information. The system has also been made to detect usage of grid resource in user's workflow in which case a grid proxy certificate will be created upon computation project submission. For the end user, it is now possible to include both local and grid resources in user's application workflow. In addition, SOMA2 program execution scheme has been extended to make use of grid middleware in actual job submission. All grid related features were implemented as optional setting and functionality is based on using user's personal X509 certificate.

7.2.1.2 Project Quarter 2

During this period, work focused on finalizing grid compatibility of SOMA2 gateway. As a result, we were able to finalize a working technical concept for submitting jobs to grid via SOMA2 that currently supports Nordugrid Arc middleware. Also, we started to prepare a new release of the SOMA2 gateway that would include features for the grid compatibility and some minor enhancements. In addition, evaluation for suitable scientific applications to be attached as part of the grid enabled SOMA2 service offering started.

7.2.1.3 Project Quarter 3

During PQ3 of the first project year, we continued to investigate suitable application candidates to be integrated in SOMA2 along with the maintenance and operations of CSC's SOMA2 service. On the programme development, no progress was made during the PQ3.

7.2.2 Current work and Outlook for 2011

During the upcoming six months, we plan to be able to release a grid-enabled version of SOMA2. Also, we intend to create a SOMA2 capsule for the chosen pilot application. This will enable to set up a grid-enabled application service, initially on a national level, in CSC's SOMA2 provided for Finnish academic researchers. This has been our target for the first project year. We will also participate in EGI User Forum 2011 in Vilnius where we are going to present SOMA2 in the portals session.

Later on, during the second project year, we will expand the service so that it will be available for other user communities as well. We will investigate possibilities to set up this service. Also, our goal will be to expand scientific applications selection in the SOMA2 service, and integrate application services from different grid entities into SOMA2. This will require a support for other middleware to be added. This should be set as an important milestone because from the end users point of view, this would make using scientific applications in different grids very easy and transparent. In addition, we intend to operate and support the SOMA2 service and we will continue to advertize the upgraded SOMA2 service to existing user communities. Also development of SOMA2 gateway will continue according to feature roadmap including possible feature requests from the user community, bug fixes and other enhancements.

7.3 CCMST

The efforts of the CCMST community are devoted to support the implementation of electronic structure and molecular dynamics codes on both HTC and HPC platforms, to facilitate the usage of related packages on the grid and to test and develop as well some tools for the assemblage of the implemented codes into efficient applications of higher level of complexity. The community makes use of both HTC and HPC platforms to develop computational tools enabling the computational chemistry, molecular and materials science community to tackle related scientific and innovative technology transfer problems. This is currently based on the efforts of gathering together a large number of members of the community and on the provision of some electronic structure packages for ab initio calculations of intra- and inter-molecular interactions, on the development and offer of some programs for quantum and classical molecular dynamics, on the design of some tools for facilitating the use of the grid, on the extension to computational chemistry codes of some libraries and tools for an efficient running of the related codes on distributed environment.

7.3.1 Work accomplished

7.3.1.1 Project Quarter 1

The main effort of the community has been devoted mainly to guaranteeing the continuity of the services previously provided by the two main VOs (COMPCHEM and GAUSSIAN) of the Computational Chemistry cluster and based on the offer of some established electronic structure packages on the distributed platform. Additionally offered in an experimental fashion was some code for quantum dynamics calculations of the reactivity of few atom systems. The testing of MPI versions of some Linear algebra routines for the grid was also started. Finally a search of centres and research laboratories to gather together for the construction of a larger CCMST community on the grid was started.

7.3.1.2 Project Quarter 2

In the second quarter, development of an improved offer of the electronic structure calculations to the wider community was developed through the assemblage of an improved version of Insilico Lab (a CYFRONET tool) in order to facilitate the usage of electronic structure calculations by a wider portion of the Molecular science community. At the same time for the molecular dynamics calculations the design of a new tool called GriF has been started by the UNIPG group in order to facilitate the study of chemical reactivity using quantum rigorus approaches for atom diatom systems. Both tools are built over other grid tools available within EGI and have a SOA oriented approach. The gathering of other research centres and laboratories has progressed by including in the community the University of Cologne and the ENEA laboratories as active members of the potential CCMST Virtual Research Community. The testing of MPI versions of some Linear algebra routines for the grid with special concern for those adopted in the quantum dynamics programs offered to the VO members of the community has progressed and reports have been made to the specific WP of EGI. Finally, the activities developed in the building of the community have been organically inserted into those of the Heavy User Board of EGI Project.

7.3.1.3 Project Quarter 3

In the third quarter the development of a strategy for the assemblage of a larger CCMST community has been the main focus of the activities. In particular the tools available within EGI for the assemblage of workflows allowing the composition of higher complexity applications have been investigated. An analysis of the use of some portals and a comparison of previous experiences carried out with the use of P-Grade has been performed. As a result it has been decided to assemble a specific e-Science environment for the CCMST community that in the meantime has strengthened the links with ENEA-GRID and MOSGRID VOs and has established new links with the large scale facility of CESCA and with the grid dissemination consortium COMETA. The testing of MPI has been extended to higher complexity codes like CHIMERE devoted to the modelling of secondary pollutants production in the atmosphere (an application of large social impact) by carrying out its profiling for parallel and distributed implementation.

7.4 Fusion

The fusion community makes use of the top computational environment with a wide range of applications, tools and developments. Several different studies and use cases can be established by means of these applications and tools. For example, the characteristics of plasma confinement can be analysed by considering the MHD equilibrium plus the Mercier and Ballooning stabilities.

As stated, there is a wide range of applications. Several of them are Monte Carlo codes that are perfectly suitable for the grid. There are also MPI applications that can be used. In order to perform these studies, these different applications must be coupled. In this coupling Kepler plays an important role in order to assure this procedure is carried out transparently to the user. 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.

7.4.1 Work accomplished

7.4.1.1 Project Quarter 1

During this quarter of year the Workflow & Schedulers support team started the coordination of tasks for the rest of the project. A survey with all the different support tools regarding the elements involved on this task was carried out, so an overall view of the of the main elements for this task was established.

A second activity concerned the preparation of a survey of the possible types of workflows required by the users and the functionalities that are needed for the implementation of these workflows that are not presently included in the workflow engine Kepler.

Finally, a lecture on “Workflows on Fusion Research” was given in the 7th International Grid and e-Science School that was held in Valencia during July 8th, 2010. The course is aimed at students and researchers who want to run applications in the GRID-CSIC, and within international initiatives EGI (European Grid Initiative), and national e-Science. In this talk the main activities in fusion research that involve complex workflows both on grid and HPC were presented.

7.4.1.2 Project Quarter 2

During PQ2, the activity was focused on establishing collaboration with other user groups interested on the actions being carried out by the activity. Presentations on the subject during the EGI Technical Forum held in Amsterdam, triggered interest from the community in the capabilities of workflow scheduling technologies. Questions and topics arose in an informal manner. It was felt that the Forum represents a great opportunity to bring together all the collaborators interested in grid technology and looking for different options to solve the challenges they have to face.

Some efforts focused on the coordination of further developments were made. The task of supporting GridWay from Kepler was designed and organized.

7.4.1.3 Project Quarter 3

During PQ3 the activity was focused on starting the work to support GridWay in Kepler based on the previous efforts done during PQ1 and PQ2. The first steps towards the support of GridWay in Kepler were taken.

One of the first tasks was to build and exploit scientifically linear workflows. These workflows do not require the use of different infrastructures to be executed. An example of this type of workflow is the execution of VMEC application (MHD equilibrium code) plus COBRA, Mercier, and analysis of the output of VMEC. VMEC is executed on one computational resource of the grid while COBRA, Mercier, and the analysis of the output is executed in another resource. This workflow uses the template actors developed to support the execution of grid applications in a transparent and reliable way.

7.4.2 Current work and Outlook for 2011

As indicated in MS602, a list of possible workflows among applications is shown in the following table.

VMEC + DKES	VMEC + ISDEP	FAFNER + ISDEP
ASTRA + TRUBA	ASTRA + TRUBA + FAFNER	EUTERPE + ISDEP
GEM + ISDEP	ASTRA + GEM	ASTRA + GEM + TRUBA

Table 3 – List of possible workflows

All the workflows have been analysed and designed in details and work needed for running each of them was evaluated in details. Basic services needed for the execution of such workflows have been deployed and setup.

Among these workflows, the Astra – Truba (4) has been built already. More workflows from those proposed in the previous list will be built and exploited in the next months. The scientists involved in the different applications are currently working on the modules which will be able to convert the output of an application into the input of another. This step is always needed by any given workflow in order to achieve an optimal communication between different applications. We will work in the next step with the workflows (1), (2), (3) in the first place.

7.5 MPI

The MPI sub-task produces numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry, 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.

To date, over 120 EGI Sites using the gLite middleware support MPI. The number of ARC and UNICORE sites has not yet been enumerated.

7.5.1 Strengths, Weaknesses, Opportunities and Threats (S.W.O.T.)

The following have been identified as the key strengths, weaknesses, opportunities and threats.

7.5.1.1 Strengths

MPI is regarded as an important tool in all scientific computing domains. The MPI standards (versions 1 and 2) and their implementation frameworks offer a well established means to enable parallel job processing. A large number of supporting tools and scientific libraries are available (open source and commercial packages). There is also a substantial body of published work showing how certain algorithms and methods fit particular workflows. This allows user communities to design and implement new MPI based applications.

7.5.1.2 Weaknesses

MPI support is dependent on numerous external factors, including batch system support and middleware support. Generic support of parallel jobs is under-developed (OpenMP, CUDA, Hadoop “map/reduce”, SmartGridRPC etc) and depends on direct interaction with the middleware providers. The range of grid-enabled MPI applications is still quite small.

7.5.1.3 Opportunities

Large SMP nodes (24 core/48 core) and GPGPUs are now available as commodity off the shelf items. The development and use of hybrid-framework applications (MPI/CUDA, MPI/OpenMP, OpenMP/CUDA) will have vastly superior runtimes to current implementations.

7.5.1.4 Threats

The range of potential site configurations demands flexible configuration and deployment tools. Middleware support for parallel jobs and, in particular, GPU support also requires early standardisation and adaptation to ensure community consensus and to avoid multiple incompatible solutions.

7.5.2 Work accomplished

7.5.2.1 Project Quarter 1

- During this quarter the MPI support team started the coordination and setup of tasks for the rest of the project. A new mailing list for internal discussion was created, the wiki with information of the task was updated and contact with the middleware developers was established.
- The "MPI Working Group Recommendations" document, initially produced under EGEE-III, was updated and finalised following dialogue between middleware developers, the MPI working group, and the user communities. The updated document clarifies the semantics of several new proposed variables that would enhance the Job Description Language (JDL). These new additions allow for greater flexibility and control of MPI jobs, however, the implementation requires changes to the gLite Workload Management System and CREAM CE. The finalized recommendation document is available for download[13].
- New training material for MPI courses was prepared and used in a training event at the Grids & e-Science 2010 course (Valencia, Spain). This material is available for download[14].
- Extension of the documentation useful to the molecular science community (COMPCHEM in particular) has been considered for offering to the users examples of MPI use at three levels of software complexity.
- The MPI support activity has continued with the support in the GGUS system with the management and resolution of any MPI related tickets during this quarter.

7.5.2.2 Project Quarter 2

- In August 2010 the MPI subtask members produced input for the milestone document “MS602 - HUC Software Roadmap”. This outlines the main objectives of the MPI subtask over the subtask’s lifetime.
- At the EGITF-2010 meeting in Amsterdam, September 2010, the MPI subtask members convened an MPI session, which was open to all interested parties – i.e. the users, VOs and site administrators. The meeting was well attended, with over 35 people present. The focus of the meeting was to relate the goals and objectives of the MPI subtask, and to introduce the support for MPI under all three middlewares. In addition, a presentation covering the measured improvements for MPI support, as measured by the Computational Chemistry VO, confirmed improvements in MPI job success rates over the previous year.

7.5.2.3 Project Quarter 3

- The MPI subtask members produced input for the EGI-InSPIE paper:

  “European Grid Infrastructure – an Integrated Sustainable Pan-European Infrastructure for Researchers in Europe (EGI-InSPIRE): The Digital Agenda for European Researchers”.
- UNIPG produced a report entitled “WORK PERFORMED BY THE RESEARCH CLUSTER OF THE UNIVERSITY OF PERUGIA: MPI on the Grid”.

---

[13] http://grid.ie/mpi/wiki/WorkingGroup?action=AttachFile&do=get&target=MPIWG-recommendation-1.0-final.doc

[14] http://indico.ific.uv.es/indico/contributionDisplay.py?contribId=15&confId=317

- CSIC created new MPI-Start documentation[15] for users and sites. A supplementary trouble-shooting guide accompanies this. MPI-START code and other developer information is also hosted at IFCA[16].
- An abstract entitled "MPI hands-on training" was submitted to the EGI-UF 2011. The training event is two hours in duration and is led by Enol Fernández del Castillo (CSIC).
- TCD joined the EGI Training Working Group and the EGI-UF 2011 Programme Committee.

7.5.3 Current work and Outlook for 2011

This year the MPI sub-task will focus on its core objectives:

- User documentation updates and improvements
- Outreach and dissemination (EGI UF, April 11th-April 15, Vilnius, Lithuania)
- User Community, NGI and site engagement, and feedback and requirements gathering
- MPI cookbook
- MPI workbenches for Computational Chemistry and Fusion Communities.

[15] http://grid.ifca.es/wiki/Middleware/MpiStart/

[16] http://devel.ifca.es/mpi-start

8 SA3 METRICS

The following table shows the SA3 metrics for the completed quarters (Project Quarters 1 – 3).

Metric ID	Metric	Task	Q1	Q2	Q3	Comments
M.SA3.1	Number of VOs deploying their own dashboard instance/view	TSA3.2.1	4	4	4	ALICE, ATLAS, CMS, LHCb
M.SA3.2	Number of users of deployed dashboard instances	TSA3.2.1	Up to 8100	Up to 8700	Up to 8600	Unique IP addresses
M.SA3.3	Number of unique users of GANGA	TSA3.2.2	839	756	692	
M.SA3.4	Number of unique users of DIANE	TSA3.2.2	17	14	18	
M.SA3.5	Number of sites using GANGA	TSA3.2.2	87	81	82	
M.SA3.6	Number of sites using DIAI	TSA3.2.2	12	9	15	# user domains
M.SA3.7	Number of users of GReIC	TSA3.2.3	75	75	~100	
M.SA3.8	Number of users of Hydra	TSA3.2.3	0	0	0	Service not yet delivered
M.SA3.9	Number of users of SOMA2	TSA3.2.4	30	33	18	Current SOMA2 service is “restricted” to CSC users
M.SA3.10	Number of users using Taverna to access EGI resources	TSA3.2.4				Work has not yet started. (See 4.2 above)
M.SA3.11	Number of users using RAS	TSA3.2.4	0	0	5	Service not yet delivered
M.SA3.12	Number of users using MD (Kepler)	TSA3.2.4	0	0	5	Service not yet delivered
M.SA3.13	Number of users using Gridway	TSA3.2.4	5	7	7	
M.SA3.14	Number of MPI support tickets	TSA3.2.5	0	0	0	
M.SA3.15	Mean time to resolve MPI support tickets	TSA3.2.5	N/A	N/A	N/A	

M.SA3.16	Number of HEP VO support tickets	TSA3.3	973	1015	929	Sum of ALICE, ATLAS, CMS and LHCb
M.SA3.17	Mean time to resolution of HEP VO support tickets	TSA3.3	246:02	182:26	241:38	HHH:MM
M.SA3.18	Number of Life Science Users of provided services	TSA3.4	8	13	14	# people in biomed technical team
M.SA3.19	Number of databases integrated and/or accessible from EGI resources.	TSA3.4	2	2	2	1 in the context of the Climate-G testbed (metadata DB) 1 for training purposes (in the context of GILDA).
M.SA3.20	Number of unique users of VisIVO	TSA3.5			15	
M.SA3.21	Number of data sets accessible from EGI resources	TSA3.6			2 (+)	

Table 4 - SA3 Project Metrics for Project Quarters 1 – 3

9 CONCLUSIONS

In this document we have presented the work and achievements of the Heavy User Communities since the start of the EGI-InSPIRE project in May 2010 together with the outlook for 2011. This document should be read in conjunction with the other deliverables of this project year which describe the Capabilities offered by these communities [D6.1] as well as the on-going work on Sustainability plans [D6.2].

Not only has a great deal been achieved during this period but also the work has successfully resulted in a number of new tools that are shared by VOs with others identified for further work in the coming year.

The VOs supported by this work are in agreement that the current focus on commonality is an important part of a strategy for the long-term, lowering the cost of on-going support and leveraging technologies shown to be useful in one VO in others.

10 REFERENCES

MS105	MS105 Quarterly Report 1:May-July 2010: https://documents.egi.eu/public/ShowDocument?docid=156
MS106	MS106 Quarterly Report 2:August-October 2010: https://documents.egi.eu/public/ShowDocument?docid=248
MS107	MS107 Quarterly Report 3:November 2010-January 2011: https://documents.egi.eu/public/ShowDocument?docid=361
MS601	MS601 – HUC Contact points and the support model: https://documents.egi.eu/document/91
MS602	MS602 – HUC Software Roadmap: https://documents.egi.eu/document/230
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
MS605	MS605 - Training and dissemination event for all shared services and the other tasks within the activity: https://documents.egi.eu/document/326
MS606	MS606 – HUC Software Roadmap: https://documents.egi.eu/public/ShowDocument?docid=310
MS607	MS607 – Hydra service deployment: https://documents.egi.eu/document/327 (to appear)
MS608	MS608 – Integration of the VisIVO server with the production infrastructure: https://documents.egi.eu/document/328 (to appear)
D6.1	D6.1 – Capabilities Offered by the Heavy User Communities: https://documents.egi.eu/public/ShowDocument?docid=154
D6.2	D6.2 – Sustainability plans for the activities of the Heavy User Communities: https://documents.egi.eu/public/ShowDocument?docid=309
D6.3	D6.3 – Annual Report on the Tools and Services of the Heavy User Communities: https://documents.egi.eu/document/312
MONARC	MONARC Technical Notes: http://www.cern.ch/MsONARC/docs/monarc_docs.html
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
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)
PHEDEX	Software agents in data and workflow management; T. Barrass et al, Computing in High Energy and Nuclear Physics (CHEP'04), Interlaken, Switzerland, 2004