EGI-InSPIRE

HUC SOFTWARE ROADMAP

EU MILESTONE: MS614

Document identifier:	EGI-MS614-final
Date:	**17/02/2012**
Activity:	**SA3**
Lead Partner:	**INFN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/746

Abstract

This document provides the fourth overview of the Roadmap for the development and deployment in the reference user communities of the software included in EGI-InSPIRE SA3. The Heavy User Communities (HUCs) who are part of the project are the primary target of the document, which is intended to give them information on the features available now and in the future, and offer the opportunity to interact with the planned developments so that they can best fit their needs.

However the document is open to other EGI-InSPIRE users and potential users, who may be interested in adopting parts of the software for their uses and in suggesting developments to this effect.

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Laura Perini, Jamie Shiers	INFN/SA3	17/02/2012
Reviewed by	Michel Drescher Marios Chatziangelou	EGI.eu NA2	03/04/2012
Approved by	**AMB & PMB**		17/04/2012

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	17/02/2012	First draft	Laura Perini/INFN Jamie Shiers/CERN
2	03/04/2012	Revised Draft after Reviewers comments	Laura Perini/INFN Jamie Shiers/CERN
3	17/04/2012	Final version	Laura Perini/INFN Jamie Shiers/CERN

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

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

VI. TERMINOLOGY

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

VII. PROJECT SUMMARY

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

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

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

The objectives of the project are:

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

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

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

VIII. EXECUTIVE SUMMARY

This report provides a snapshot of the status and planning (for year 2012) of the services and tools developed and supported for the needs of the Heavy User Communities (HUCs): High Energy Physics, Life Sciences, Astronomy and Astrophysics, and Earth Sciences. They have provided information both on the software that is currently specific for each of them and for the software that is already of interest for more than a single community.

This last category consists of the same services and tools considered in the previous versions of this Milestone: the Dashboards, GANGA and related applications, the HYDRA and GrelC services, the Kepler, Gridway, SOMA2 workflow schedulers, and the enabling of MPI applications, which receives important contributions also by the Computational Chemistry HUC.

This report is a checkpoint for the goals of SA3, recalled briefly in the points below.

- To transition the services and tools from the communities that have already adopted DCIs, to the point where their services are part of the general service infrastructure provided through EGI or are sustained by other means – either within their own community or through external software providers (e.g. middleware projects such as EMI).
- To use the experiences obtained by these early adopting communities in integrating new data sources, tools, and services in order to improve the experience for all user communities.
- To ensure that all the user communities supported by EGI should experience no disruption as they move on away from their current e-Infrastructure provider.

The report and planning of the services and tools, sketched in this report, show relevant developments and widespread usage, mainly still concentrated in one or two communities. The relevant effort going towards making the use of these services and tools easier, while at the same time increasing their functionalities as required by the users, promises well for their more general use within SA3 and also outside of it.

Table of Contents

\

1 INTRODUCTION

This document reflects the present status in the elaboration of a full roadmap (the HUC Software Roadmap is a Milestone due periodically in SA3): most of the software developments are still concentrated towards a single community; however, much effort is devoted to the documentation and to the easiness of use of the different products, both of them being necessary conditions for attracting new user communities. For the different software products the planning included in this document covers at least the main features foreseen in the next year, in some cases providing well defined internal milestones and initial indications for longer term developments.

New versions of this document have been produced every 6 months, starting from MS602 [MS602] in Project Month 4. The communication between the different communities has started and is growing – this is a visible and concrete deliverable of the project where it is clear that the funding model and goals of the project are succeeding in motivating common tools and services, even if at the architectural but sometimes also at the implementation and deployment level; the planning for their software of potentially more general interest, has been exposed to the other communities of heavy users and to the general users communities.

2 ROADMAP FOR THE SA3 SHARED SOFTWARE SERVICES AND TOOLS

2.1 Dashboards

2.1.1 Overview

The Experiment Dashboard [DASHBOARD] system provides monitoring of the WLCG infrastructure from the perspective of the LHC experiments and covers the full range of their computing activities, such as data transfer, job processing, and site commissioning. In contrast to many other monitoring systems, the Experiment Dashboard is not coupled to any particular middleware, workload management, or data management systems. It is shared by several LHC virtual organizations and works transparently across various middleware platforms.

The Experiment Dashboard proved to be an essential component of the Grid monitoring infrastructure. The LHC experiments heavily rely on the Experiment Dashboard for the distributed operations, for monitoring user computing activities, and for providing user support. The Experiment Dashboard servers are accessed by several thousand distinct users per month. Among the Dashboard applications which are extensively used by the LHC community are Site Status Board, Job monitoring applications, DDM (Distributed Data Management) Dashboard.

Experiment Dashboard applications adhere to a set of core development principles: common technology and implementation, loose-coupling to data sources, sharing of monitoring data, and user involvement in the development process. The aims of these principles are to reduce development and maintenance overhead, to allow applications to be easily adapted for use by multiple VOs, to enable reuse of monitoring data, and to ensure that applications meet user requirements. The Experiment Dashboard Framework provides a common foundation for the development of the monitoring applications and facilitates development and maintenance tasks.

2.1.2 Sustainability

The sustainability of the system is ensured by the described development principles and by the collaborative nature of the project. The core development and support team at CERN collaborates with the developers from the LHC experiments. There are also contributions from the institutes participating in the LHC project from Russia, Taiwan, India, and the UK. The LHC experiments also take part in operating the Dashboard services, for example as a part of ATLAS computing shift procedures.

During 2011, substantial progress was made concerning the redesign of the Dashboard applications, achieving complete decoupling of the user interfaces from the data repository through the use of new technologies such as AJAX. New versions of the Site Status Board, of the ATLAS DDM Dashboard and Site Usability Monitor, as well as of the wide range of the job monitoring applications were deployed in production. Many of the job monitoring applications were developed in a common hBrowser visualization framework, which aided the development process and allowed the integration of common user interfaces with various information sources. The migration of the Dashboard cluster consisting of 50 nodes to virtual machines was accomplished. The functionality of the Dashboard cluster management system was improved. This includes development of the web interface which provides information about status of the nodes belonging to the Dashboard cluster and applications hosted by these nodes. All those changes contributed to the sustainability task through decreasing maintenance and support effort.

In 2012, priority will be given to the common applications shared by multiple VOs. Among them is the Global WLCG transfer monitor. The Global WLCG transfer monitor aims to perform common tasks carried out by the LHC virtual organizations in order to monitor their data transfers on the WLCG infrastructure, namely aggregation of the WLCG transfer monitoring statistics and exposing monitoring data via user interfaces and APIs. The first prototype of the system was developed and deployed for validation in 2011. During 2012, the functionality of the system will be extended based on feedback of the user community.

Another important direction of work will be the improvement of the monitoring tools for end users, task monitoring for distributed analysis being one of them. A new version of the task monitoring application based on the hBrowser framework will be released in 2012. All Dashboard data repositories will be migrated to new hardware and ORACLE version 11g; this will improve performance of the Dashboard applications. The functionality of the Site Status Board application will be extended in order to provide better automation of the operations of the distributed infrastructure, which in its turn will address the sustainability task. It is foreseen to be integrated into the SSB framework functionality of the CMS site readiness tool and to enable alarm notification. Apart from the development work described above, the plans for 2012 also include maintenance and support of the Dashboard cluster, services, and software.

2.1.3 Roadmap

In 2012, priority will be given to the common applications shared by multiple VOs. Among them is the Global WLCG transfer monitor. The Global WLCG transfer monitor aims to perform common tasks carried out by the LHC virtual organizations in order to monitor their data transfers on the WLCG infrastructure, namely aggregation of the WLCG transfer monitoring statistics and exposing monitoring data via user interfaces and APIs. The first prototype of the system was developed and deployed for validation in 2011. During 2012, the functionality of the system will be extended based on feedback of the user community. Another important direction of work will be the improvement of the monitoring tools for end users, task monitoring for distributed analysis being one of them. A new version of the task monitoring application based on the hBrowser framework will be released in 2012. All Dashboard data repositories will be migrated to new hardware and ORACLE version 11g; this will improve performance of the Dashboard applications. The functionality of the Site Status Board application will be extended in order to provide better automation of the operations of the distributed infrastructure, which in its turn will address the sustainability task. It is foreseen to be integrated into the SSB framework functionality of the CMS site readiness tool and to enable alarm notification. Apart from the development work described above, the plans for 2012 also include maintenance and support of the Dashboard cluster, services, and software.

2.2 Applications

2.2.1 Ganga

2.2.1.1 Overview

Ganga is an end-user tool for creating and managing computational tasks. Its mature Python codebase provides for a stable yet extensible framework, which is consistently used by over 400 unique end-users each month. It remains a popular environment for running grid analysis jobs within the LHCb and ATLAS experiments, and is known to be used by a further 10 user communities across a range of scientific disciplines.

Development of the Ganga core, the central, generic code around which community-specific applications are built, is led by the CERN IT department and driven by the evolution of requirements of Heavy User Communities (HUCs). A major core development milestone of 2011 was the incorporation of a 'Prepare' method into the generic framework. This functionality was previously only available within the GangaATLAS package, but has since been implemented for a range of applications. Current development within the core is focused on providing generic output-handling routines which allow users to pass their output through a range of post-processing steps. For example, output files can now be compressed before being shipped back to the user or, alternatively, sent to remote storage technologies.

Future developments will concentrate on facilitating the adoption of Ganga as an analysis platform. For example, though the current installation mechanism is relatively simple, inclusion of Ganga within a recognised Linux distribution, such as EPEL, the Extra Packages for Enterprise Linux maintained by a Fedora special interests group, would not only provide a familiar installation procedure for new users, but also reinforce the fact that Ganga is a mature and established tool. Prior to inclusion, a review will be held to identify any code changes needed to ensure Ganga is compatible with the default EPEL Python version. At the same time external dependencies will be reviewed, since retirement or evolution of certain Ganga functionality may have rendered them obsolete.

2.2.1.2 Sustainability

Long-term sustainability of the Ganga project is aided by its modular architecture. The Ganga core contains a common software framework that is shared by all user communities and their applications, in addition to providing the software process infrastructure (i.e. packaging and distribution tools, internal test and release tools and services). This commonality ensures that the core packages can be supported with reduced manpower whilst serving multiple communities. Indeed, during 2011 Ganga was adopted by a new physics community who continue to develop their own application that may be incorporated into the official Ganga codebase in the future.

The ease with which Ganga can be integrated with community-specific services is a result of the design architecture employed. A well documented[1] plug-in approach decouples the core components from the experiment specific framework and insulates Ganga somewhat from changes in those frameworks, which should be transparent to the core (or at least incorporated with minimal effort).

Ganga supports several job submission backends, including the most widely supported and utilised job submission services in EGI, i.e. the Cream CE and the gLite WMS (beside Condor and various other batch systems). This widens the appeal of Ganga, and ensures the project is resilient to future policy changes within a community. Similarly, the flexible framework allows additional backend interfaces to be developed with minimal effort, future integration with Cloud computing infrastructures being one potential example.

In addition to the above technical design approaches, the community-based nature of Ganga helps ensure its sustainability, with key development themes being shared amongst a number of developers at different institutes, thereby mitigating the impact of personnel loss to the overall project. Furthermore, the development team actively seeks and encourages new user communities to adopt Ganga, develop their own applications, and engage in regular development meetings, workshops, and discussion forums. Taking such an open approach allows external parties to contribute to the Ganga codebase, with the benefit of boosting the total person-power available to the project.

---

[1] https://twiki.cern.ch/twiki/bin/view/ArdaGrid/GangaIndex#Developers

2.3 Services

2.3.1 GrelC

2.3.1.1 Introduction

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

During the EGI-InSPIRE project, the GRelC middleware will be enhanced to support the HUC communities with a new set of functionalities available as a web application both through the GRelC Portal and the DashboardDB interface.

2.3.1.2 Future Plans for the next 12 months

The support in terms of management, monitoring, and control of the GRelC services provided through the DashboardDB will be further extended and improved. The DashboardDB will provide new (service-specific) views about the GRelC deployment and the status of the registered services. A strong dissemination about the DashboardDB registry will be carried out over the next period to register the GRelC service instances already deployed in the EGI grid as well as the grid-database resources they offer. Moreover, the questionnaire prepared during Y1 will be improved, extended, and disseminated among the HUC to attract new users and make the existing ones aware of the new grid-database registry tool.As for the gLite 3.2 release, the integration of the GRelC middleware into the Italian Grid Initiative (IGI) release will be completed. A roadmap towards the *EMI Release* will be defined as well.

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

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

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

2.3.1.3 Longer term developments

The GRelC roadmap towards the European Middleware Initiative release may be considered an important point for the next period. In any case, it should be considered a long term task, since it will need continuous support if it is to be implemented over the next years.

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

during Y3. This study will help in better defining a longer term roadmap related to this action.

Finally, the longer term development plan foresees the design and implementation of new high level tools able to:

(i)   Ease the access and management of a *single GRelC service*;

(ii)   Provide high level functionalities involving a *set of GRelC services* ("collective layer") to run data integration, analysis, and mining over a grid environment.

2.3.1.4   Sustainability plan

Sustainability can be directly (applying for different sources of funding) or indirectly (by working on technical/scientific aspects that enable the involvement in new proposals) addressed.

In the former case, several proposals related to national and international calls have been submitted over the last years. Specific calls addressing HUC have been considered. For instance one of them related to the Italian LifeWatch Virtual Laboratory recently got the approval. Another grant, started in the second half of 2011 (three years term), concerns the extension of the GRelC software to include On-Line Analytical Processing (OLAP) functionalities for climate change data management. Other proposals including some extensions of the GRelC service to efficiently manage climate metadata have been also submitted.

In the latter case, indirect aspects that make the participation and the involvement into new proposals possible and that are strongly taken into account in the GRelC roadmap are: a larger adoption of the software, a real and effective support in terms of functionalities for HUC issues and use cases, a strong level of dissemination and visibility, an easy integration into existing systems, the scientific value (which can be evaluated for instance in terms of publications on relevant journals), the ability to address common needs across different communities defining some exploitation patterns, the re-usability of the software, its robustness and performance. Visibility is a major point and it has been addressed updating the website and the wiki, giving oral presentations and tutorials, writing new papers and organizing sessions and workshops in ICT (e.g. International Conference on Computational Science - ICCS2011) and domain-related conferences (e.g. European Geosciences Union – EGU2011), as well documented in the Quarterly Reports.

All of these aspects (both direct and indirect) play a significant role in addressing sustainability and, till now, they have been taken into strong consideration with regard to the GRelC middleware. Of course, part of the work over the next year will be related to strengthening these aspects, trying to better address them.

2.3.2   HYDRA

2.3.2.1   Description

Hydra [R2] is a file encryption/decryption tool developed by EMI to enable the protection of sensitive files stored on grid storage resources. The service is composed by a distributed encryption keystore (hence its name), and client command line tools that can (i) upload/fetch keys to/from the keystore and (ii) encrypt/decrypt data files using these keys. For maximal protection, the keys are split into pieces and are distributed over the keystore, so that they are partially redundant (e.g. 2 out of 3 key pieces may be needed to reconstruct a complete key) and incomplete (e.g. at least 2 pieces are needed to reconstruct a complete key) following the Shamir's secret sharing algorithm [R1].

2.3.2.2   Provision status

An experimental Hydra service has been successfully deployed since the end of Q4 (April 2011), first on a 32bits gLite version 3.1 and then on a 64bits gLite version 3.2 server. It is usable for test purposes. A 3 servers-based Hydra keystore was planed to be deployed on servers from the CNRS I3S and CREATIS laboratories, and from the HealthGrid association. However, the dissolution of the HealthGrid association after judicial decision prevented the completion of this deployment plan and so the service is still available as a single-server keystore today. Two other difficulties have been identified as far as delivering a production service using the Hydra keystore is concerned. Firstly, many production sites are misconfigured, having deployed older version of Hydra in a former gLite release and sometimes exposing hydra tags that are not valid anymore or older version of the hydra client. Secondly, the current version of the hydra client is developed as part of the EMI middleware and its dependencies are incompatible with the gLite 3.2 release, deployed on the production infrastructure. The service delivered today is, therefore, only a test service, mostly used for the validation of the functionality delivered, including:

- o   Encryption key registration and removal.
- o   Encryption key access control.
- o   File encryption / decryption from User Interfaces and Worker Nodes.

The following table summarizes the service provision status and delivery plan:

Task description	Current status	Plans
Hydra keystore server provisioning	Latest official installation and configuration procedure available [R3] for gLite 3.0. Unofficial customized procedure applies to gLite 3.1/SL4 [R4] and gLite 3.2 (undocumented). A prototype keystore was deployed on a 64 bits gLite 3.2 machine.	The provision of proper Hydra packages for the EGI infrastructure depends on EMI. It will probably not be available until the first release of EMI is deployed on the production infrastructure.
Hardware provisioning	Single keystore available for testing. HealthGrid was supposed to host a keystore. However it was dissolved in January 2012.	Supporting NGIs are currently being requested to provision a server for the hydra keystore.
Sites configuration	Many computing sites have incorrect configuration resulting from the deployment of former Hydra software package releases.	Life Sciences supporting sites are being contacted to clean-up their configuration.
Make Hydra client package available to all Worker Nodes by pre-installing it on the shared area of each site accessible to the LS HUC VOs,	The current EMI packages cannot be deployed on the production infrastructure.	It is unlikely that EMI will provide gLite-compatible packages. There is an ongoing attempt to adapt EMI packages to the current production

and publishing the corresponding runtime environment tag in the BDII for those sites.		middleware. Whether that fails will depend upon EMI's release installation date on the EGI infrastructure. An ultimate option would be to ask supporting sites to deploy the previous gLite-compatible release of the hydra client package on their Worker Nodes.
Publish the Hydra service in the BDII so that users can access it through the gLite Service Discovery [R5] simply based on the service unique id.	The registration of the Hydra service shall be done only when the production service is installed. Until then, the experimental service will be using the local-file based configuration.	The procedure is being discussed between EMI and BDII experts as to the most appropriate way to register Hydra in the gLite Service Discovery.
Monitoring: set up probes to (1) monitor the service from an end-user point of view (encryption/decryption commands succeed), and (2) make sure each keystore server is up and running.	Not started.	New probes to be developed within the Biomed Nagios box.

2.3.2.3 Sustainability

To deliver a production service, Hydra depends upon:

- Servers provision for hydra keystores;
- Standard client packages installation on all EGI computing resources;
- Service monitoring and maintenance.

In order to ensure a sustainable service and to avoid difficulties encountered with the HealthGrid association, NGIs will be requested to host the hydra server store. The Life Science Grid Community (LSGC) technical team will handle the service client installation and monitoring for more sustainability. An attempt will be made to produce client packages compatible with the current middleware release and to deploy them. If this proves impossible, it will be needed to wait for the first deployment of the EMI release on the EGI infrastructure. The service monitoring procedure will be integrated in the Nagios box-based monitoring service offered by the LSGC technical team.

2.4 Workflow and Schedulers

2.4.1 Kepler and GridWay

As already indicated in MS606, Kepler [KEPLER] is a software application for the analysis and modelling of scientific data. Kepler allows scientists to create their own executable scientific workflows by simply dragging and dropping components onto a workflow creation area and

connecting the components to construct a specific data flow; GridWay [GRIDWAY] is a Metascheduler that automatically performs all the **submission steps** and that also provides the runtime mechanisms needed for dynamically adapting the application execution.

The work of these activities has been focused on building a set of different use cases that can be used as templates by any other scientific community. Up until now, the developed workflows have been using the capabilities of gLite or UNICORE, with the functionalities provided by the standard software of these middlewares. Since some of the applications (and users) were already using Gridway, we decided to also provide support for it from within the Kepler environment.

Before getting started with this work, we needed to have a use case which could benefit from this kind of support. The use case we have just created, FAFNER+ISDEP, will allow us to test the functionality of the integration of GridWay within Kepler. This development has been delayed up to this point because of the necessity of the use case. Thus, the plan for the next twelve months includes integrating GridWay into Kepler and testing this integration with the use case already developed and/or future use cases.

GridWay provides functionalities that have been already developed in the Serpens suite for Kepler. because of the characteristics of the sofware supported (WMS, LB,...). For example, with GridWay, more actions can be delegated into the metascheduler since this software deals with most of the casuistry regarding job management: send the job, retrieve the results, deal with the information about the status of the infrastructure, etc..

By supporting GridWay, we are improving the sustainability plans of Kepler. This development will provide enhanced capabilities plus new options for the users. GridWay supports different middlewares, hiding those middlewares from the user or from any piece of sotware on top of GridWay. Hence, more computing infrastructures (middleware) could be used from within Kepler.

As it has been stated in previous reports, activities have been focused on establishing collaboration with other user groups which have shown interest in our work. Also, the dissemination activities accompanying various meetings and conferences made the work visible to other communities. We have established contact with the compchem users, and are in the phase of helping out with the development of the first use cases using Kepler (with the Serpens suite). We also plan to approach other user communities. As for the short term dissemination activities, we plan to perform hands-on training sessions during events like the EGI User Forum.

2.4.2 SOMA2

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

During PQ6 and PQ7, CSC has maintained and operated CSC's SOMA2 service. In the recent SOMA2 development work, the main outcome has been the integration of scientific application Autodock 4 into SOMA2 so that it makes use of features that enable job execution in a grid. Furthermore, at the end of PQ7, SOMA2 version 1.4 Aluminium was released. This release contains most of the EGI-InSPIRE related work to date, and we have continued to investigate on how to setup a SOMA2 service which could be provided to other user communities, not just those from CSC.

2.4.2.1 Outlook for future work

Based on the achievements so far, we will setup a SOMA2 service which provides access to scientific applications by making use of grid resources. This way the current SOMA2 service will be extended from a national level to other user communities as well. For this purpose CSC is going to set up a separate pilot service. We will also operate and support the existing SOMA2 services. In addition, we will advertise the grid enabled SOMA2 service to different user communities.

For the longer term, our goal is to expand the selection of scientific applications in the SOMA2 service, and to integrate application services from different grids hosted by different virtual organizations into SOMA2. This should be set as an important milestone because from the end users point of view, this would make using scientific applications in different grids very easy and transparent. We will continue to maintain and operate the SOMA2 service and seek possible scientific applications to be added as part of the service and during EGI-InSpire we will support communities/VOs using SOMA2 services.. Furthermore, the development of the SOMA2 gateway will continue according to the feature roadmap, including possible feature requests from the user community, bug fixes, and other enhancements.

2.4.2.2 Sustainability

SOMA2 has been in production use at CSC for many years. In addition, the SOMA2 source code was released as a downloadable package under the GPL open source license in 2007. To date, CSC has been able to establish a solid user base for SOMA2 at a national level. Our motivation for developing grid support in SOMA2 was to extend the current CSC centralised national use to an international level. We expect that the open source release of the SOMA2 and pilot service for EGI user communities (and the EGI collaboration in general) will lead to an increased number of SOMA2 deployments. This in turn will help to extend the developer community outside CSC, hence providing sustainability for SOMA2 open source software.

Currently CSC develops and maintains the SOMA2 service and this activity is expected to continue for the national effort after the EGI-InSPIRE project, providing a minimum level of continuation for SOMA2 beyond the EGI Inspire project. The goal is that some EGI VO's will also take part in supporting the service in the future.

In general, SOMA2 has a modular design in which scientific applications are integrated in the system with scientific application descriptions called SOMA2 capsules (plugins). This provides the possibility to establish not only a software developer community but also a community for the service providers where specialists could exchange information about application integration, which typically is not only a technical issue but also requires in depth knowledge about the applications to be used. For example, a repository for sharing SOMA2 capsules could facilitate this.

Finally, sustainability of SOMA2, method of deployments (centralized gateway vs. local installations), continuation and extension of the EGI pilot and all community aspects will depend on the level which SOMA2 is adopted within the EGI communities. We are using this opportunity to promote SOMA2 within EGI and its applicable VO's.

2.5 MPI

2.5.1 Overview

The MPI sub-task shall focus on a number of core objectives over its 36 month period (PY1 -PY3).

Many of these objectives are iterative, often requiring updates or fine-tuning. Other objectives, including participation in the EGI Community and Technical forums, shall be repeated at regular intervals. The core sub-task objectives 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.

This sub-task has produced numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry and Fusion communities. These products include parallel implementations of Linear Algebra routines (produced by UNIPG). UNIPG has been working on the parallelization of the CHIMERE application. The detailed analysis[CHIMERE] has identified the extent to which the application will achieve the highest level of parallelization possible, and thus the greatest efficiency. This model is structured around a task-farming model, in which some worker processes share the work by taking charge of part of the domain (subdomain) under the control of the master that sends and receives the relevant data, and that writes the results. The application can read initial input data based on the output of a previous run for a given timeframe. The model has several advantages:

- By chaining the results of many shorter runs, this allows simulations to run over much longer timeframes, and thus overcomes issues with CPU time-bound usage limitations at most resource centres.
- Can also be used as a form of application check-pointing.
- Allows greater granularity/resolution.

The CHIMERE parallelization model can be reused in many other scientific domains including, among others, the Earth Science and Astronomy & Astrophysics communities. UNIPG has also implemented numerous advanced parallel codes for GPGPU. This code has been successfully implemented on two different GPGPU architectures. In addition, UNIPG shall continue exploration of a parallel implementation of CHIMERE and the GPGPU work in PY3.

Improvements to the MPI documentation[2] were made in PY2. The EGI Wiki now provides the definitive source of information for MPI support. Documentation on the ARC and UNICORE MPI support still require extra attention. The current release of the UMD gLite-MPI 1.2.0 package was released in Project Month 21. The release of the UMD glite-WMS and the glite-MPI products are expected to contain significant middleware changes required to ensure the correct support for many MPI job types. For example, these include:

- User defined allocation of processes/nodes

2

https://wiki.egi.eu/wiki/MPI_User_Guide

- Direct Support for OpenMP codes.

Furthermore, end-user support for both MVAPICH (an MPI implementation with advanced Infinband networking) and support for generic parallel job support is a high-priority in PY3.

This work is under the leadership of CSIC, and is aimed at steering MPI and Parallel jobs technical requirements.

Both UNIPG and TCD are working on exploiting GPGPUs. The UNIPG work is already at an advanced stage: An entire job flow that enables the Local Resource Management System (LRMS) to discriminate the GPGPU resources requests, through Glue Schema parameters, has been defined in order to allocate, in a dynamic fashion, the required resources on a Cloud-like infrastructure. This solution helps overcome some weakness in many job schedulers, which do not currently support GPGPU resources. TCD uses StratusLab – an open-source cloud distribution that allows grid and non-grid resource centres to offer and to exploit an "Infrastructure as a Service" cloud; this is a work-in-progress, and significant results are expected in PQ8. TCD is investigating potential GPGPU Information System support.

As part of User Community engagement effort, the MPI team will regularly survey Virtual Organisations, Users and Site administrators for critical feedback. This will also act as a means to gather information about current deficits and future requirements. The newly founded NA2 driven MPI Virtual Team (MPI-VT) is expected to lead to a concentrated drive to weed out and fix problems observed with some large production-level MPI jobs. This shall create a greater demand on the SA3 MPI support team, led by CSIC and TCD.

The predominant job scheduler "MAUI" needs to be upgraded at affected MPI sites in PY3, due to the fact that a node allocation bug/problem has been discovered in the standard distribution. EMI does not provide support for MAUI, and the software license may impose significant software distribution restrictions. We are currently working around this problem by recommending that sites install a patched version from an unofficial repository. Longer-term support requires that this issue be resolved at the highest level.

2.5.2 Sustainability

The MPI task has not yet fully tackled the issue of sustainability. The MPI community itself has two mature and major open-source framework implantations: MPICH-2 and OpenMPI. These are both well supported and are expected to be maintained and further developed in the medium term (five years). Both are also heavily used by the High Performance Computing community. The two frameworks are also actively developed, and produce several new releases per year. The latest overview of the OpenMPI roadmap can be found in the "State of the Union" document[OpenMPI-SOTU].

 In the medium term, both MPICH and OpenMPI are working on contibuting and implementing the MPI-3 draft standard[MPI-3]. In particular, the latest development version of MPICH2, version 1.5, implements some of the draft standard (in particular, non-blocking collectives). In addition, both are attempting to integrate CUDA support to allow MPI applications to copy data to and from available GPGPUs. This is currently available in MPICH2 version 1.4.1p1, and in OpenMPI 1.5. These are not enabled by default, and require Resource Centre customisations. CUDA support depends on the

OpenMPI "hwloc" sub-project.

Methods to widely support generic parallel jobs, GPGPU, and cloud integration may be potential research projects for further exploitation on distributed computing infrastructures.

3 ROADMAP FOR THE COMMUNITY SPECIFIC SOFTWARE

3.1 Services for the HEP Virtual Research Community

Services for High Energy Physics are covered in detail in MS610 [MS610].

In the following sections a brief intermediate report and planning is provided covering the next 6 months, the schedule on which this document is updated and replaced.

3.1.1 HammerCloud

3.1.1.1 Overview

[HammerCloud] is a Grid site testing service developed around Ganga. HammerCloud uses frequent short jobs to validate a site's availability and functionality, and also delivers on-demand stress tests to aid in site commissioning or general benchmarking. Ganga was developed with the ATLAS LHC experiment and that collaboration continues to use the service actively. The CMS and LHCb LHC experiments also have HammerCloud plug-ins developed.

Achievements in 2011 include:

- Completion of the deployment of HammerCloud for CMS and LHCb.
- Deployment of the auto-exclusion system for the ATLAS experiment that automatically excludes sites from the distributed analysis systems. Improvements in reliability have been measured, reducing the failure rate of Grid jobs up to 50%.[3]
- Extension of functional testing to production queues for ATLAS.

Plans for 2012 include:

- Migration of the ATLAS instance to the new version, merging development with the CMS and LHCb instance.
- Development of features for new database backends based on NoSQL engines (such as CouchDB) for the use case of high demand storage (mainly metrics and job results).
- Extension of the auto-exclusion feature that has been tested in production with ATLAS and proved effective, to the CMS and LHCb experiments.
- Decommission of Job Robot for CMS, switching to HammerCloud as the main Grid testing service for CMS along with ATLAS.
- Optimization and generalization of submission mechanisms to reduce load.
- Deeper integration with LHCb DIRAC and further deployment for LHC*b* Grid testing.
- Integration of LHCb testing with the LHCbDIRAC resource status system.
- Working on a standard benchmark (or benchmarking suite) to evaluate the storage elements (e.g. standard set of ROOT test jobs and datasets).
- Investigation of error and performance correlations across VOs (at multi-VO sites).
- Development of a multi-VO testing interface for users to easily test all experiment workflows.
- HC installation simplification and packaging for other-VO usage.

3.1.1.2 Sustainability

The HC project is led by an EGI-INSPIRE-funded FTE, and developments and operations are complemented by at least 1 partial FTE from each user VO (ATLAS has 2 extra support persons

---

[3] D. van der Ster. "Automatic site blacklisting for analysis". WLCG Collaboration Workshop, DESY. 11 July 2011.

(partial)). The exact plan for the project leadership past the end of EGI-INSPIRE is not known at this time. However, it is expected that the continued success of the HammerCloud project does not depend on any one partner.

3.1.2 CRAB and CRAB Analysis Server

3.1.2.1 Overview

The CMS Remote Analysis Builder (CRAB) [CRAB] was the first analysis tool in CMS to aid users in configuring CMS applications for distributed use [CMSDA], by discovering the location of remote datasets and submitting jobs to the Grid infrastructure. CRAB has progressed from a limited initial prototype nearly 5 years ago to a fully validated system that was critical to the production of published physics results. CMS currently observes more than 400 unique users, submitting ~900K CRAB jobs per week, with close to 1000 individuals per month. The CMS Computing Technical Design Report (CTDR) [CTDR] estimated roughly 100k Grid submissions per day. During the second half of 2010 the job submissions routinely exceeded the estimate by 40-50% and CRAB coped well with the increased load.

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

The main items of development for 2011 are:

- To rewrite all the monitoring implementation, building centralized service, based on CouchDB, able to collect monitoring documents from the distributed agents.
- To implement advanced functionality like user data publication, support for the user output merging.
- To evolve the current client allowing its usage as a python library. The aim is to enable the possibility to be used by external bot/services like HammerCloud, which will be the first case we will support.
- To rewrite the RESTFul based APIs using a better modular approach, which means a set of common libraries among various Web Services of the CMS Workload and Data management. The new APIs are going to be rewritten mostly to address the security aspects.
- To evaluate and eventually integrate and use frontier on worker node to cache UserFileCache
- To perform a long scale test, including the whole stack which compose the CRAB3 system, namely: Client, REST Interface, UserFileCache, WorkQueue and distributed Agents.

3.1.2.2 Sustainability

CRAB is a joint project among various institutes, mostly CERN IT-ES, FNAL and INFN. If any of these teams decrease the effort spent on the project, although there will be an impact on CRAB, it will not affect the sustainability issue. As an example, CERN is soon to celebrate its 60[th] anniversary and has recently outlined plans for the next 40 years. FNAL and INFN also measure their life in decades with funding expected well into the future. Their lifetime can thus be expected to greatly exceed the useful lifetime of any one software or middleware product.

3.1.3 Data Management

3.1.3.1 ATLAS Distributed Data Management

The ATLAS Distributed Data Management (DDM) is the project built on top of the WLCG middleware and is responsible for the replication, access, and bookkeeping of the multi-Petabyte ATLAS data across the grid, while enforcing the policies defined in the ATLAS Computing Model. The current ATLAS DDM software is now in a mature state and the present work is focused on maintenance and support operations.

Inside the ATLAS Distributed Computing community, there is an on-going discussion about the proposal of the CERN-PH-ADP group of developing a new DDM system (the Rucio project [RUCIO]) to solve the current shortcomings and scalability issues in the Central Catalogues. The details about the future of the project and the implications in other groups are yet to be negotiated and are unclear at this point.

3.1.3.2 CMS GridData Popularity Frame-work and Automatic Site Cleaning

Following first experience with real LHC data taking, production, and analysis, it has become clear that some of the assumptions behind the current model are not strictly valid. For example, the "pre-placement" of data from Tier1 to Tier2 sites for analysis has shown that a large fraction of the data is never accessed. CMS is working on more dynamic data placement strategies that can result in better resource utilization (potentially lower network traffic as only the needed data is transferred and improved storage management through smaller numbers of data copies).

For this purpose, CMS has deployed its popularity framework that holds the information about which data is used for analysis at the different Grid sites according to the official CMS workload management infrastructure (CRAB). In the last quarter we implemented a system for the dynamic cleaning of unused replicas when the different physics groups exceed their allocations on sites.

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

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

3.1.4 Persistency Frame-work

3.1.4.1 Overview

Persistency and detector conditions cover two essential areas for HEP data processing and analysis, namely the handling of the event data and detector conditions data of the LHC experiments. Event data, that record the signals left in the detectors by the particles generated in the LHC beam collisions, are generally stored on files, while conditions data, that record the experimental conditions

---

(like voltages and temperatures) at the time the event data were collected, are commonly stored using relational database systems. In three of the LHC experiments (ATLAS, CMS and LHCb), some or all of these types of data are stored and accessed inside data processing jobs using one or more of the three software packages developed by the Persistency Framework [PERS-FRAME] project: CORAL [CORAL], COOL [COOL] and POOL [POOL].

The three Persistency Framework packages specialise in different areas of data persistency. COOL, which provides specific software components to handle the time variation and versioning of conditions data, is the only package that is strongly HEP-specific. POOL is a generic hybrid technology store for C++ objects and object collections, using a mixture of streaming and relational technologies. CORAL, a generic abstraction layer with an SQL-free API for accessing data stored using relational database technologies, is the package that could potentially be most interesting for other communities than HEP (when compared to COOL, which deals with a HEP-specific data model, or POOL, whose maintenance is being discontinued also for HEP, as discussed below); for this reason CORAL was in 2011 and will again in 2012 be presented at the EGI User Forum. CORAL supports several relational database backends and deployment models, including local access to SQLite files, direct client access to Oracle and MySQL servers, as well as read-only access to Oracle through a middle tier server and a caching proxy using the FroNTier/Squid [FRONTIER] and the CORAL server/proxy [CORAL-SERVER] technologies.

The section below discussed the ongoing efforts for clarifying the support model for CORAL, COOL and Frontier, with the goal of achieving sustainability for their usage in HEP. Apart from these discussions at the management level, on a more technical level the plans for 2012 include: user support, service operation and software maintenance (as discussed above); the development of new functionalities, such as the capability of the CORAL software to restore the database connection and session after a network glitch (expanding on previous work in 2011); R&D to prototype an improved monitoring infrastructure for CORAL, and more particularly for the CoralServer and CoralServerProxy components in the context of their usage for the ATLAS HLT system.

3.1.4.2 Sustainability

The Persistency Framework software has been developed over several years (since 2003 for POOL, since 2004 for CORAL and COOL) through the well established collaboration of developers from the LHC experiments with a team in the CERN IT department (now partly funded by the EGI-InSPIRE project), which has also ensured the overall project coordination. Within the common project, the personnel pledged by ATLAS, CMS and LHCb, coming from a large number of institutes in several countries, have contributed to the development and continue to support the components used by their experiment. The common project, in particular the effort from the IT Department (and from EGI-InSPIRE through it) only deals with components that concern (or, initially, that show the potential to concern) more than one experiment: in the past, individual experiments have already taken up the sole responsibility of components that have been proven to be relevant to that single experiment. This is in line with the more general focus on common solutions as a strategy for the sustainability of the HEP community. The usage of the three packages is periodically reviewed with the relevant stakeholders in the IT Department and the experiments. In particular, the sustainability of the support model for POOL, CORAL and COOL, taking into account the upcoming termination of funding from EGI-Inspire, has been discussed within the context of the WLCG Technical Evaluation Groups in PQ4 2011.

For what concerns POOL, these discussions have already successfully converged on an agreed sustainable support model. LHCb has already stopped using POOL, after replacing it by a new

software layer, inspired from POOL but maintained internally, which is also able to read existing data stored in POOL format; as a consequence, LHCb no longer needs support from POOL through EGI and its software has no pending dependencies. ATLAS will continue to use POOL for as long as the 2012 production version of the ATLAS software is actively used, but it will no longer use it or need POOL support through EGI for the 2013 production version, where a custom software package derived from POOL will be built and maintained by ATLAS as part of their internal software. A development version of the 2013 ATLAS software release already exists; this will eventually become the production version of the ATLAS software, by the end of 2012 or beginning of 2013.

For what concerns CORAL and COOL, a sustainable support model for the future is still being discussed with all its users (ATLAS and LHCb for both, CMS only for CORAL) and other relevant stakeholders. While the software is by now mature in its development cycle, a large development and support effort (approximately 3 FTEs) is still required for user support, service operation and maintenance tasks. Software releases are prepared as requested by one or more of the LHC experiments, leading to one release per month on average: this is generally motivated either by functionality enhancements or bug fixes in the Persistency Framework packages, or by upgrades in the versions of the "external" dependencies (ROOT, Boost, Python, Oracle...). In this respect, an important step towards sustainability was achieved in 2011 by transferring the full responsibility for CORAL and COOL release build and validation, from the EGI-funded IT team to the team in the PH department that takes care for software infrastructure issues common to all LHC experiments. Service operation incidents and user support requests normally result in bug fixes or functional enhancements in the CORAL and COOL code, but may often end up in need of a more global analysis involving other software packages (such as Oracle, the grid middleware, or ROOT). In particular, understanding service operation issues for CORAL-based applications almost always requires detailed troubleshooting on the Oracle database servers where the data is stored; these tasks are performed by the Persistency Framework team in collaboration with the relevant database administrator team (funded by CERN IT). As an example, a major effort and success for COOL in 2011 was the validation of query performance in the move from Oracle 10g to 11g servers, which required the detailed analysis of some issues that were eventually identified as a bug in the Oracle server software, fixed since that time.

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

Apart from the clarifications of the support model for CORAL, COOL and Frontier at the management level, on a more technical level the plans for 2012 include: user support, service operation and software maintenance (as discussed above); the development of new functionalities, such as the capability of the CORAL software to restore the database connection and session after a network glitch (expanding on previous work in 2011); R&D to prototype an improved monitoring infrastructure for CORAL, and more particularly for the CoralServer and CoralServerProxy components in the context of their usage for the ATLAS HLT system.

3.1.5 DIRAC

3.1.5.1 Overview

The DIRAC system provides a complete solution for using the distributed computing resources of the LHCb experiment. DIRAC has a complete framework for data production and analysis, including workload management, data management, monitoring, and accounting (more details have been given in document [MS610]). One of its most important components is the Data Management System (DMS), whose support in EGI-InSPIRE project started in October 2010. During the first year of the project, significant progress has been made in the improvement and development of the system, mainly in the area of data consistency between storage elements (SEs) and file catalogues and the accounting of space usage at Grid sites.

3.1.5.2 Plans for the Future

The main lines of development for the next year are the following:

- Consolidation of the service for data consistency checks between file catalogues and SEs, both in the development and operational aspects. This service, developed during last year, will be improved on the basis of the experience gathered with its usage during the past months. In particular, procedures should be established for notifying the VO's data management team about the files which were found in an inconsistent state, and decide about their removal. This will streamline the task of consistency checks and will enable the data manager to reduce the time needed to fulfill the task.
- Support and possible improvement of the system for the accounting of storage resources usage developed and put in production during the past year. This system, already extensively used by members of the collaboration, could undergo some improvement, taking advantage of new features of the LHCb DIRAC DMS that are planned to be implemented during next months.
- Adapt when necessary the DIRAC DMS to change in the data management middleware (FTS, SRM, LFC etc.).
- Continue the development of the popularity service, just started during last PQ. Such a service should provide metrics to asset the data-sets popularity and provide a ranking of the most popular data-sets (i.e. data most frequently accessed by users). The final goal is to use the information about data popularity to implement a dynamic data placement model, where the number of replicas of a given data-set is related to its popularity. This would considerably help the VO's data managers to optimize the usage of storage resources on the Grid and would automate many operations that are currently done manually and thus time expansive.
- Provide general support for LHCb computing operations on the grid, both for production and private users' activity.

3.1.5.3 Sustainability

The LHCb DIRAC project has made considerable progress towards a sustainable future by trying to improve the model for developing software, certifying and making releases. For what concerns software development, a modular approach is always privileged, in order to avoid any duplication of code and making the program more easily maintainable. For the process of certification and releases, a more systematic process has been recently defined where a more careful testing is required from developers. This has the positive effect of reducing the time required for certification

and for producing new releases. For next year, the objective is to consolidate these practices which have been recently introduced.

3.2 Services for the Life Science Virtual Research Community

3.2.1 Description

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. As described in MS611 [MS611], the Life Sciences HUC contributes to the LSGC effort to organize the community and deliver services.

3.2.2 Delivery status

In this context, several services to the HUC users have been provisioned. Details can be found on the LSGC wiki[4]:

- Mailing lists addressing each affiliated VO and the whole VRC have been set up and are updated daily to ensure communication within the community.
- Web gadgets listing Life Sciences applications and community requirements posted to the RT systems set up by User Communities Support Team have been added to the LSGC wiki.
- A HUC support service is delivered. A technical team of expert users has been set up. It addresses the difficulties reported by users on the VRC mailing lists or through GGUS. Bi-monthly phone conferences are scheduled and shifts are organized to ensure that there is always a team on duty tackling the problems. See technical team wiki for details[5].
- The technical team pro-actively monitors the infrastructure health at a VO level, to identify the problems occurring. The French NGI deploys a dedicated Nagios server[6] for that purpose. New probes to monitor all VO SEs, WMSs and CEs were developed.
- On-line reporting tools easing the monitoring of SEs space management have been added to the technical team wiki page.

The LSGC recently faced the dissolution of the HealthGrid association upon juridical decision. HealthGrid provided human support and services hosting to the LSGC. As a consequence, the redundant VOMS server previously deployed is not available anymore and the plan for the provision of a Hydra keystore server has been interrupted as detailed in the section above.

Previous plans to deploy a redundant LFC server are currently on hold due to technical incompatibilities encountered (the biomed LFC server is using an Oracle database back-end and its replication would require finding a site providing an Oracle licence for that purpose; a migration to MySQL back-ends is not recommended for performance reasons).

---

4 LSGC wiki, http://wiki.healthgrid.org/LSVRC:Index

5 Biomed technical team wiki, http://wiki.healthgrid.org/Biomed-Shifts:Index

6 Biomed Nagios server, https://grid04.lal.in2p3.fr/nagios

3.2.3 Sustainability

Despite the unexpected dissolution of the HealthGrid association, the LSGC organization proved to be resilient. In the future, technical services hosting requests will be made out to the NGIs to improve LSGC services sustainability.

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

3.3 Services for the Astronomy & Astrophysics Virtual Research Community

3.3.1 Overview

Activities carried out by the A&A community in the period from September 2011 to February 2012 have been focused on the following topics:

- Coordination of the A&A community with particular reference to the mid and long term sustainability plan;
- Tools and Services for visualization: VisIVO;
- HPC, parallel programming, and GPU computing;
- Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure.

3.3.2 Roadmap

3.3.2.1 Coordination of the A&A Community

The coordination of activities within the A&A community concerning the usage of DCIs for both small-scale and large-scale projects, in particular the ESFRI ones, i.e. SKA (Square Kilometer Array) and CTA (Cherenkov Telescope Array) has been intensified. A coordinating workshop[7] took place at the Astronomical Observatory of Paris on November 7th 2011, where the major astronomical projects and research areas were represented.

During this workshop the current status of the A&A HUC was discussed. The most important topics raised from the discussion highlighted the need of having multiple astronomical VRCs each of them aggregated around a big project or Institution well representative of one of the major A&A research areas. A unique catch-all VRC, in fact, has proven to be not suitable to represent the whole community and poses strong limitations for what concerns the requirements gathering process and the preparation and implementation of a sustainability plan. All participants who attended the workshop agreed that the major effort within the A&A HUC should be focused on the identification of big A&A projects and Institutions for which the adoption of DCIs could be beneficial, so they could be

---

7

http://twiki.oats.inaf.it/twiki/bin/view/AstroVRC/AstroVRCWorkshop

good candidates to lead new astronomical VRCs.

3.3.2.2 VisIVO

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

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

Further developments of VisIVO include: a) the implementation of a VisIVO web portal for gLite; b) the production of a MPI- compatible version of VisIVO for gLite; c) the production of a CUDA-compatible version of VisIVO for gLite. All these activities are already in progress.

3.3.2.3 Grid and HPC

Activities related to Grid and HPC are carried out in close coordination with EGI and in Italy with IGI, the Italian NGI. One of the goals of this activity is the introduction of small-size HPC resources in Grid. This implies, on one hand, the installation and configuration of HPC clusters (based on low latency/high throughput networks, HPC libraries and tools, modules and compilers) and , on the other hand, making the grid middleware aware of these resources. In this context, A&A applications able to take advantage of such small-size HPC resources include FLY (a cosmological code developed at INAF-OACt) and Gadget + Flash, the most common cosmological codes in Astrophysics. These applications can be used to test and validate the new resources and the Grid middleware.

This activity may also have a significant impact on other disciplines that need HPC and MPI in order to efficiently run their applications.

GPUs are emerging as important computing resources in Astronomy; in fact, they can be successfully used to effectively carry out data reduction and analysis. The option of using GPU computing resources offered by Grid sites to make visualization processing on VisIVO is currently under evaluation.

A joint activity has been recently discussed and planned within IGI with the involvement of its major communities; this activity aims at introducing new HPC-oriented tools as well as new resources. In the next period therefore the activity of task TSA3.5 for what concerns HPC will continue in close coordination with all other communities of IGI in Italy.

3.3.2.4 Access to Databases and interoperability with the VObs

Access to databases from DCIs and interoperability with the Virtual Observatory data infrastructure remains one of the hot topics in astronomy. One of the main related activities mainly deals with the identification of use-cases and test-beds (both applications and complex workflows) that require simultaneous access to astronomical data and computing resources. The data infrastructure of reference in astronomy is the Virtual Observatory; end users should be able to access astronomical data in Grid through the VObs standards[8] and launch computational tasks on DCIs. In this context two key issues to tackle concern the SSO (Single Sign On) mechanism able to grant access to computing and data resources through a single authentication phase (users do not need to authenticate themselves multiple times) and a tool/service to access astronomical databases federated in the

---

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

VObs from DCIs. GRelC is one of the tools under evaluation to verify its ability to meet the most important requirements of the A&A community. The evaluation of GrelC will be one of the main activities of the next period together with the provision of requirements and use-cases in close coordination with the VObs.

3.3.3 Sustainability

VisIVO is a visualization tool whose early design and implementation has been carried out focusing on the needs of astronomical applications; as a consequence of its late evolution, however, it is now a general purpose visualization tool, able to meet the requirements of many communities and of their applications, therefore its future sustainability should not represent a problem.

Similarly, topics such as Grid-HPC and the interoperability between DCIs and databases impact many communities and EGI itself is active in finding adequate solutions for them. Therefore there are not expected sustainability issues in this case.

As highlighted in the above section dedicated to the coordination of the A&A community, a strong effort is currently in progress to identify A&A projects and Institutions that might be good candidates to lead well focused VRCs. If the attempt of pushing the coordination of the A&A community in this direction will be successful, the preparation and implementation of a sustainability plan should be easier given that the maintenance of new developed tools and services would be ensured by Institutes and projects leading the well focused VRCs.

3.4 Services for the Earth Science Virtual Research Community

3.4.1 Progress and plans

The VRC maintains and develops interfaces to data outside of the EGI Infrastructure in Earth Science data centers. The data is mandatory for simulations running in the EGI infrastructure. One of the largest metadata providers of satellite data is the GENESI Infrastructure from GENESI-DR and GENESI-DEC (FP7 projects led by ESA). During the project a client to discover and download files from GENESI was built and is still maintained. This client automatically creates the necessary JDL files for a user application including a shell script, which downloads the selected datasets from the sites that are part of GENESI. To avoid conflicts with duplicates, data is only stored on EGI resources during runtime. All required management is done by the client. Initiated by requirements from ES users, a portal solution with a flexible web GUI based on GI-cat web services developed by ESSI-Lab (http://essi-lab.eu/gi-cat) is now operational and maintained. This service extends the access to other providers using different protocols.

The access from the EGI to the Earth System Grid is not possible without credentials. The team -- IPSL and the ES VRC in EGI -- works on Earth System Grid (ESG) interoperability with considerable progress. A testing scenario is based on an application from IPSL. This application uses climate model data stored in the ESG belonging to the fifth phase of the Climate Model Inter Comparison Project (CMIP5). The application is in preparation and a credential test bed addressing the today's different authorization methods is installed.

The maintenance of the interface to access GENESI Data is continuing. A continuous adaption to ongoing changes at the GENESI search interface and its connected data repositories sites is also

necessary. The interface is constantly updated to recent changes and adaptations in the backend system. Especially further integration with the GI-cat web service and the possibility to download data from a variety of data centers directly to a computing or storage element, besides GENESI sites, will be done in the next month.

The IPS-ES VO team will go on to develop and install a demonstrator with test bed accepting a chain of different credentials, a transfer tool for CMIP5 data (ESG -> EGI) and an application. The transfer tool will facilitate bulk-oriented access to ESG CMIP data. It will support features for users to get data from a list of models and includes multi-threaded downloading and incremental downloading through history with caching mechanism to limit the stress on ESG metadata servers. The credential test bed in place will be used for further testing of different authentication schemes on EGI and ESG. The development of these features has started and is ongoing in the next quarters.

3.4.2 Sustainability and cross-community aspects

As the developed software is rather focused and consists of atomic applications developed for specific use-cases of the Earth Science community, we fear that the prospective gains for other communities are rather limited. Nevertheless, the inquiries and technical developments may be of assistance, or even serve as groundwork for other communities. Although the specific technical conditions for the application areas, such as the data handling the tools deal with, discern seriously throughout the different communities, common requirements, issues, and principles could be identified. The underlying basics are however not part of this task, that is more practice-oriented. As the tools are self-contained, we foresee that they can survive as open source projects.

The investigation on authentication schemes, more generally, bridging infrastructure security functions, are of general interest. The development is aligned and communicated to both the ESGF as well as EGI. Members of the task have participated in identity management workshops and plan to continue to engage in further developments. We hope that this issue will be solved in a larger scale than the ES task of SA3, such as in the infrastructure projects, as they can be helpful in numerous situations.

4 CONCLUSIONS

All the Heavy Users Communities included in the SA3 work-package have contributed to this report on the HUC Software Roadmap with information on the status and planning of their services and tools. New versions of this document have been produced every 6 months, starting from [MS602] in Project Month 4 and [MS606] in Project Month 10. The communication between the different communities is growing – this is a visible and concrete deliverable of the project where it is clear that the funding model and goals of the project are succeeding in motivating common tools and services, even if at the architectural but sometimes also at the implementation and deployment level; the planning for their software of potentially more general interest, has been exposed to the other communities of heavy users and to the general users communities. As described in detail in the Deliverable D6.2 on Sustainability plans for the HUCs [D6.2], a variety of common areas between the HUCs have already been identified, with solutions being adopted by multiple VOs and considered also by others. Concrete examples include the HammerCloud Stress Testing System for Distributed Analysis [HC], now adopted by ATLAS, CMS and LHCb as well as the more recent work on dynamic data placement and caching (aka Data Popularity). Users wishing to know more of the capabilities of these services should consult D6.4 "Capabilities Offered by the HUCs to Other Communities" [D6.4], as well as MS609 "HUC Contact Points and Support Model" [MS609]. The achievements of the work package in the first year of the project are detailed in D6.3 "Annual Report on the HUC Tools and Services" [D6.3].

As defined by the mandate of this work package, the feedback from the different kind of users, the sites and the software developers will continue to be an important next step towards the long term planning of these services and tools for sustainability.

5 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
MS608	MS608 – Integration of the VisIVO server with the production infrastructure: https://documents.egi.eu/document/328
MS609	MS609 – https://documents.egi.eu/public/ShowDocument?docid=419
MS610	MS610 – Services for High Energy Physics: https://documents.egi.eu/secure/ShowDocument?docid=540
MS611	MS611 – Services for the Life Science Community: https://documents.egi.eu/secure/ShowDocument?docid=683
MS612	MS612 – HUC Software Roadmap: https://documents.egi.eu/public/ShowDocument?docid=684
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
D6.4	D6.4 - https://documents.egi.eu/secure/ShowDocument?docid=472
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)
GREIC	http://www.grelc.unile.it/
HYDRA	https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS
R 1	A. Shamir. "How to share a secret" in Communications of the ACM (CACM), vol. 22, pages 612-613, 1979.
R 2	Hydra service overview. https://twiki.cern.ch/twiki/bin/view/EGEE/DMEDS
R 3	Hydra service installation. http://glite.web.cern.ch/glite/packages/R3.0/R20060502/doc/installation_guide_3.0-2.html#_Toc135537608
R 4	Hydra test set-up installation for gLite 3.1/SL4 http://euindia.ictp.it/stock-analysis-application/hydra-test-set-up-installation
R 5	Service Discovery User Guide: https://edms.cern.ch/document/578147/1
KEPLER	https://kepler-project.org/
GRIDWAY	http://www.Gridway.org/doku.php?id=start
SOMA2	http://www.csc.fi/soma Kinnunen, T., Nyrönen, T., Lehtovuori, P., SOMA2 - Open Source Framework for Molecular Modelling Workflows, *Chemistry Central Journal*, 2(Suppl 1):P4 (**2008**). Lehtovuori, P. Nyrönen, T., SOMA - Workflow for Small Molecule Property Calculations on a Multiplatform Computing Grid, *J. Chem. Inf. Model.*, 46(2) (**2006**) 620-625.
MPI	http://egee-uig.web.cern.ch/egee-uig/beta_pages/MPIJobs.html
CHIMERE	https://wiki.egi.eu/w/images/d/df/Chimerereport.pdf
OpenMPI-SOTU	http://www.open-mpi.org/papers/sc-2011/Open-MPI-SC11-BOF-1up.pdf
MPI-3	http://meetings.mpi-forum.org/draft_standard/mpi3.0_draft_1.pdf
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
PERS-FRAME	https://twiki.cern.ch/twiki/bin/view/Persistency
LSGC-WIKI	LSGC wiki, http://wiki.healthGrid.org/LSVRC:Index
RUCIO	Rucio: Conceptual Model; M. Barisits et al, ATLAS internal note reference: ATL-COM-SOFT-2011-030, http://cdsweb.cern.ch/record/1386633