EGI-InSPIRE

HUC SOFTWARE ROADMAP

EU MILESTONE: MS620

Document identifier:	EGI-MS620
Date:	**18/09/2012**
Activity:	**SA3**
Lead Partner:	**INFN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/750

Abstract

This document provides the fifth and final 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 (potential) users, who may be interested in adopting parts of the software for their uses and in suggesting developments to this effect.

Information on the Capabilities Offered by the HUCs to Other Communities can be found in deliverables D6.1 and D6.4. Information on the long-term sustainability of the tools and their support can be found in D6.3, D6.6 and D6.9 (to appear).

I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Laura Perini, Jamie Shiers	INFN/SA3	18/09/2012
Reviewed By	Karolis Eigelis, Mark Santcross		9/11/2012
Approved By	AMB & PMB		21/11/2012

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	18/09/2012	First draft	Laura Perini/INFN Jamie Shiers/CERN
2	09/11/2012	Response to reviewers' comments	SA3

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 the remaining SA3 months until end April 2013) 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 effort in 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.

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 common solutions for monitoring of the computing activities of the LHC experiments on the WLCG infrastructure. The Dashboard applications are widely used for the WLCG operations, for monitoring of analysis tasks by the LHC physicists and user support teams, for prompt detection of problems and inefficiencies of different natures, for data mining of the collected monitoring information in order to organize LHC computing activities in the most optimal way. The Experiment Dashboard is not coupled with any of the workload management system or data management system of the LHC experiments which allows it to work transparently across various VOs. Since the LHC experiments use several different middleware platforms, an important design principle used for all Dashboard applications is to find solutions which can be easily adapted for a variety of the middleware flavors or Grid services. 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, enable reuse of monitoring data, and ensure that applications meet user requirements. The DASHBOARD framework provides a common foundation for the development of the monitoring applications and facilitates development and maintenance tasks.

The current implementation of Experiment Dashboard requires ORACLE as the backend solution as it is the default supported storage backend at CERN. Adapting the workload management and data management systems of new VOs to report to the Experiment Dashboard is a prerequisite for job processing and data transfer monitoring. For monitoring of sites and services with a new user community, the Site Status Board can be deployed, but the adopting VO must develop publishers to generate monitoring metrics.

An alternative approach is the EGI Introductory Package that offers a simple but complete solution for running and monitoring of computing tasks on the grid. The Package is designed for small and medium user communities who want to start using grid technologies without minimal initial overhead. It is easy for a community to customize or scale-up the system as their needs grow. The lightweight generic dashboard version included in the EGI Introductory Package uses mySQL as a backend and monitors the Ganga and Diane job submission systems.

The role of the Experiment Dashboard is constantly increasing, in particular with the start of the LHC data taking. In the beginning of July 2012 CERN announced that the LHC

experiments observed particle consistent with long-sought Higgs boson. As has been mentioned several times during the CERN Higgs seminar, and in the following press conference, the computing has been an integral part of this major achievement. Professor Heuer the CERN Director General had a summary statement that listed the 3 essential tools "the accelerator – the experiments – Grid computing". Being an important component of the LHC Grid computing, the Experiment Dashboard provided required functionality for the monitoring of the LHC computing activities and contributed to the success of the CERN scientific program. ATLAS presentations at the CERN Higgs seminar and at the year's major particle physics conference, ICHEP2012 in Melbourne included graphical plots from the Experiment Dashboard demonstrating the ATLAS job processing activity on the Grid infrastructure.

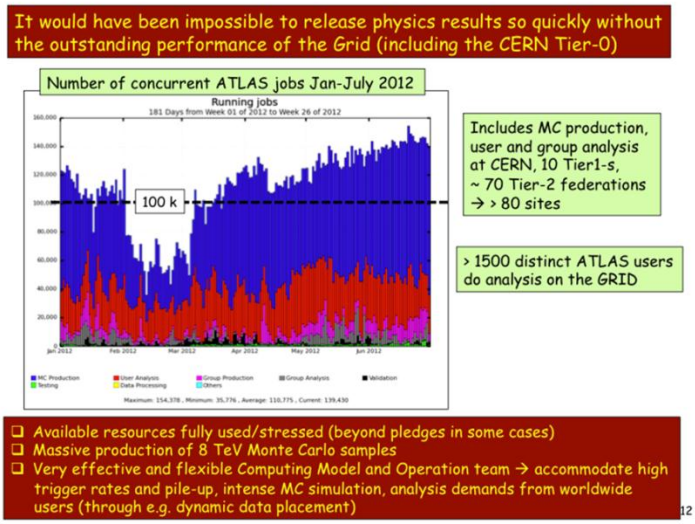

Figure 1 Slide from the ATLAS presentation at CERN Higgs seminar which contains Experiment Dashboard job processing plot. Several thousand distinct users access the Experiment Dashboard services every month and the number of users is steadily growing. The functionality of the existing applications is being extended and the new applications are being developed.

2.1.2 Sustainability

The user community takes an active role in the development of Dashboard applications and in operations of the Dashboard services. Along with the development principles described above, involvement of the user community in the development and maintenance tasks contributes to the sustainability of the Experiment Dashboard system.

During 2012 substantial progress was made in all monitoring areas: monitoring of the data processing, monitoring of the data transfer and data access, monitoring of the distributed sites and services. Priority was given to applications shared by multiple VOs.

The new system, WLCG Transfer Dashboard, was deployed in production in early summer 2012. The system aims to provide a cross-VO and cross-technology view of data transfers performed on the WLCG infrastructure. In contrast to the existing Data management systems of the WLCG experiments which work in the scope of a single VO, the WLCG Transfer Dashboard should monitor all WLCG transfer activities and should provide a global view

through a single entry point. The first version which is currently in production covers transfers which are performed using File Transfer Service (FTS). Any Virtual Organization (also outside the LHC scope) which uses FTS for data transfer can monitor its transfers using the WLCG Transfer Dashboard. Currently there are more than 15 VOs whose transfers are monitored via the WLCG Transfer Dashboard.

Another new application which was deployed in production in the beginning of 2012 is the Site Usability Monitor (SUM). Site Availability Monitor (SAM) had been completely redesigned during last couple of years and now it is based on the Nagios system. SUM represents a new portal for visualizing of the SAM test results and site availability and reliability based on these results and on various VO-specific profiles. Though the application has been initially developed for LHC VOs, who already widely use it in their everyday operations, it can be also deployed for other VOs, under the condition that they define VO-specific profiles and provide topology information in an agreed format.

One more application which provides monitoring of distributed sites and services is the Site Status Board. The functionality of this application was extended following the requests of the LHC VOs. Usage of this application contributes to the sustainability of the LHC computing operations. For example, ATLAS plans to use SSB for auto-exclusion of sites which do not satisfy the VO requirements. Such auto-exclusion will decrease the overall job processing failure rate and the need of manual interventions.

New versions of the job monitoring *interactive* and *historical* views and of the task monitoring for analysis and production had been deployed in production for ATLAS and CMS. All these applications are developed in the common h-browse framework.

In the beginning of 2012 all Dashboard repositories migrated to ORACLE 11g. This substantially improved the performance and stability of Dashboard applications. For example performance of the job monitoring applications improved 5 to 10 times.

The development plans for the next seven months will again focus on the generic applications which can be shared by multiple VOs. One of the highest priority tasks is to enable monitoring of the transfers of the xrootd federations which are being increasingly used by the LHC experiments. The xrootd transfers will be monitored via the WLCG Transfer Dashboard which currently covers only FTS transfers. The task monitoring applications will be extended in order to enable job killing and job resubmission using the Task Monitoring UI. SUM and SSB will be further improved in order to increase effectiveness of the WLCG operations and to reduce effort required to perform everyday operational tasks.

2.2 Applications

2.2.1 Ganga

2.2.1.1 Overview

Ganga is an end-user tool for creating and managing computational tasks. It remains a popular environment for running grid analysis jobs (400+ users per month) within the LHCb and ATLAS experiments, and continues to attract new communities. During the last year Ganga was adopted by four new experiments (SuperB, BES-III, SNO+ and T2K) who either use it as their end-user analysis platform or, in the case of SNO+ and T2K, as a 'production job' driver.

Recent Ganga development included a significant refactoring to decouple Gaudi framework components from the GangaLHCb application, and create a set of experiment-neutral Gaudi classes. As a result, communities wishing to adopt Ganga as their analysis tool (and who already utilise the Gaudi framework) can exploit these generic classes to create their own custom Ganga-based tools.

2.2.1.2 Sustainability

Long-term sustainability of the Ganga project is aided both by its modular architecture and also by the fact that the user base involved in Ganga development continues to expand.

Development of the Ganga core codebase (the central, generic code around which community-specific applications are built) and maintenance of the packaging, testing and software distribution tools and services is led by the CERN IT department and driven by the evolution of requirements of Heavy User Communities (HUCs).

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). Recent experience[2,3] demonstrates that this approach works, with new communities having developed their own custom applications, guided by the core Ganga development team.

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.

2.3 Services

2.3.1 GreIC

Introduction

The GRelC service is a grid database interface aiming at providing access and management functionalities related to relational and non-relational databases in a grid environment. Along with user community support, the main goal of this task, in the context of the EGI-InSPIRE project, is to design and implement a new system, named DashboardDB, which

(i) monitors the grid-database service instances and
(ii) publishes the grid-databases available in the EGI production grid exploiting a social-oriented approach.

[1] https://twiki.cern.ch/twiki/bin/view/ArdaGrid/GangaIndex#Developers
[2] https://twiki.cern.ch/twiki/bin/view/ArdaGrid/SuperBDevelopment
[3] http://indico.cern.ch/contributionDisplay.py?contribId=35&sessionId=8&confId=149557

At the end of Y2 a major release of the DashboardDB system has been delivered.
It consists of two parts:

- the *DashboardDB registry*;
- the *DashboardDB global monitoring*.

The **DashboardDB registry** aims at providing a web access interface to the database resources available in the EGI production grid. Such a "registry" complements the functionalities provided by the EGI Application Database and supports the users using social network features, e.g. a message board on the DashboardDB platform (similar to shoutboxes found on public community websites), create discussion groups, rate existing resources, etc. The users can publish their own data sources and discover new ones already deployed through the search & discovery functionality. The **DashboardDB global monitoring** allows the monitoring of the GRelC instances deployed in the EGI grid. Such a monitoring part complements the existing monitoring tools already available in the EGI environment providing a different perspective more focused on "grid-database management" services. Both the DashboardDB registry and global monitoring have been designed as self-consistent components (they can be easily exported and embedded in other web contexts like web gadgets) and are now available through "permalinks".

Plans for the next 7 months

The DashboardDB has been released at the end of Y2 and over the next months it will be updated and populated both in terms of grid-database services and grid-database resources. To achieve this key objective, a strong dissemination/communication plan has been defined in close collaboration with NA2 representatives before this summer. This plan identifies and schedules the main actions useful to reach a wider user community and it is now in its implementation phase till the end of the Y3. In parallel, the production release will be improved and fixed as new bugs will be discovered.
The future work over the next 7 months, will focus on the following four main activities:

1. extending/fixing the functionalities provided by the DashboardDB,
2. informing the users about this new tool publishing 'Scientific Databases',
3. creating an IGI compliant repository (ig_GRelC) and defining a plan towards an EMI compliant release of the GRelC middleware,
4. supporting the end users communities.

In the following, a detailed plan, with actions, expected deadlines and preliminary results is described.

Concerning point (1) the **DashboardDB registry** will provide additional "social oriented" features (most of them are now under test) like:

1.1) creating multiple discussions associated to each grid-database (rather than just a single one);

1.2) including multimedia (images/videos/urls) content (rather than only text) into users' posts and adding like/dislike and reply-to-post functionalities;

whereas, the **DashboardDB monitoring** one will include additional support regarding:

1.3) the monitoring view at the "service-level", providing service availability, round trip time, authorized VOs for each grid-DB service and number of published grid-DB

resources for each grid-DB service over the time.

All of the actions concerning point (1) are almost completed and will be available online by the end of September. A preliminary demonstration will be available in Prague during the EGI-TF2012.

Regarding point (2) the dissemination/communication plan (jointly defined in June with NA2) is going to be addressed and finalized by the end of October/November, through the following actions:

2.1) publicize the 2 gadgets on the EGI gadgets website and create a new entry under 'Support Services' of the EGI website related to 'Scientific databases' (this action is already ongoing);

2.2) implement a communication plan including the preparation of a blog post, newsletter item and/or leaflet to attract new users (both end-users and grid-dabase service providers);

2.3) work with the NILs on identifying and linking scientific databases to EGI with GRelC;

2.4) contact other grid-database service providers to include into the DashboardDB registry/monitoring system, additional ‘Scientific Databases’ managed through different grid-database interfaces.

Regarding point (3) there are basically two actions:

3.1) validation of the IGI repository ig_GRelC. The repository is already in place and right now the yaim configuration part is under test. This activity should be completed by mid October. A bug discovered in July on the Globus external libraries RPM has slighly delayed this task;

3.2) definition of a plan towards the European Middleware Initiative (EMI). Since 2007, the GRelC software is a production-level grid-database service compliant with gLite. In this regard, a release plan towards the compatibility of this software with the EMI distribution is now under evaluation and will be defined by the end of PQ10. The deadline to provide preliminary outcomes on the porting activity is the end of Y3.

Regarding point (4) the support to the HUC will be performed till the end of Y3 through the following actions:

4.1) provide support to the end users needing help to implement their use cases. This is the case of the Biology group at the University of Salento that is implementing a couple of Life Sciences use cases (see previous Quarterly Reports) and the Euro-Mediterranean Centre on Climate Change (CMCC) that is using GRelC as grid metadata service. Moreover a new use case jointly defined with CMCC scientists and exploiting GRelC as a grid data interface to manage multidimensional (NetCDF) datasets has been defined and will be implemented till the end of Y3;

4.2) update GRelC and DashboardDB websites with new material and information;

4.3) post some presentations about GRelC on the EGI Training MarketPlace;

4.4) present this activity in domain oriented conferences (like the European Geosciences Union) to identify new use cases and attract new users as done in the past years.

Finally, some of the users that have used/are using the GRelC software are being contacted to give an update about their experience with this service and to check their availability to

publish their services and grid-databases on the DashboardDB system. This process (already ongoing) will last until the end of Y3. In this regard, as a preliminay result, two sites (one in Catania - INFN-CATANIA - and another one in Naples - GRISU-SPACI-NAPOLI), are going respectively to update/install the gLite 3.2 version of GRelC, publishing these new resources into the DashboardDB system.

Longer term developments

The GRelC roadmap towards the European Middleware Initiative release may be considered a long term action. It will need continuous support even after Y3.

Moreover, 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.

Finally, the longer term plan foresees the design and implementation of:

i. specialized and highly efficient domain-oriented support (at the "service" level) to carry out data analysis and mining (this activity has been proposed into an Italian grant started in the second half of 2011, see next section);
ii. high level functionalities involving a *set of GRelC services* (at the "collective" layer) to run dataflow tasks, data federation/integration activities over a grid/cloud environment (dataflow topics are part of a FP7 biodiversity proposal under preparation, whereas data integration/federation tasks will be designed and implemented into another biodiversity project, at the Italian level, which recently got the approval, see next section);
iii. new high level systems (more "community oriented", like the DashboardDB) to attract new users starting more from the dissemination of the available database resources and trying to build a community around them (this was the case of the activity proposed for GRelC in the EGI-Inspire proposal).

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 needs (in particular in the Environment and Biodiversity domains) have been considered. For instance one of them (related to the Italian LifeWatch Virtual Laboratory) recently got the approval. It foresees database integration/federation use cases in the biodiversity domain. Another one in the FP7 framework is under preparation and will include topics related to grid/cloud-based database management for correlation analysis across different domains. Another grant, started in the second half of 2011 (three years term) and concerns the extension of the GRelC software to include On-Line Analytical Processing (OLAP) functionalities for climate change data management. Another project named TESSA started in January 2012, will last until the end of 2014 and it is going to use the GRelC service as part of the data management infrastructure to manage the metadata of two distributed oceanographic data archives (related to the Mediterranean sea) for Situational Sea Awareness. The GRelC service instances that will be deployed in the two TESSA sites (Naples and Lecce) will be included into the DashboardDB monitoring and new users from this project (from 15 to 25) are expected to register to the

DashboardDB system. Concerning the GRelC service, further extensions (oriented to the management of "timeseries for oceanograghic data") will be developed during the project.

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: dissemination and visibility, an easy integration of the software into existing systems, a scientific background (in terms of papers/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 through tutorials, updated news, material and documentation on the project websites as well as defining and working on a strong dissemination/communication plan in close collaboration with NA2.

Both **direct** and **indirect** aspects play a significant role in addressing sustainability and, till now, they have been taken into strong consideration into the GRelC roadmap. Of course, part of the work over the next years will be related to strengthening these aspects, trying to further address them.

2.3.2 Hydra

2.3.2.1 Description

Hydra [Hydra] is a file encryption/decryption tool developed by EMI to enable the protection of sensitive files stored on storage resources. The service is composed of a distributed encryption key store (hence its name), and client command line tools that can (i) upload/fetch keys to/from the key store and (ii) encrypt/decrypt data files using these keys. For maximal protection, the keys are split into pieces and are distributed over the key store, so that they are partially redundant (e.g. only 2 out of 3 key pieces are 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 [Shamir].

2.3.2.2 Provision status

The experimental Hydra service deployed within the first year of the EGI-InSPIRE project on a single gLite release 3.1 UI has been migrated to two gLite release 3.2 servers from the CNRS I3S and CREATIS laboratories, with the help of CERN experts from the EMI project, see [MS607].

A new partner was identified to provision the third Hydra key store, in order to replace the server initially provisioned by the HealthGrid association, after it was dissolved. This partner, CNRS "Institut Pluridisciplinaire Hubert Curien" (IPHC), has taken this opportunity to start updating Quator scripts to deploy the current non-official release of the Hydra key store, that should be part of a later EMI2 delivery (not scheduled yet).

The Hydra service relies on the fact that the Hydra client software be (i) installed on all sites where Worker Nodes may be required to access the Hydra service (presumably all sites accessible to the LS HUC VOs), or (ii) installed and published by means of runtime environment tags on those sites that wish to support the service. However, a survey has revealed that lots of production sites were misconfigured, not having deployed the Hydra client, having deployed an older version of the Hydra client, or publishing Hydra tags that are

not consistent with the deployed client if any. Consequently, during this period, a negotiation was led with each site publishing Hydra tags, or providing Hydra client without tags, to clear off the situation.

Still, additional concerns hamper the deployment of the service in production: on one hand the Hydra client currently deployed on the production infrastructure is available only with gLite 3.2, which security support will terminate by November 2012. On the other hand, there is no planned date for the delivery of the Hydra client developed as part of the EMI middleware, and first tests show that its dependencies are incompatible with the gLite 3.2 release. Hence, along with resource providers choosing to migrate their Worker Nodes to EMI before the Hydra client is released, we can expect the number of appropriate WN supporting Hydra to decrease.

Therefore, in such an unclear situation, the service delivered today remains a test service mostly used for the validation of the functionality delivered and the testing of the deployment procedures.

2.3.2.3 Future plans

A close follow-up of EMI release plans is organized to ensure that the client be integrated in future EMI release and deployed as part of the standard distribution. This work is bound to the frequency of EMI release schedules though.

2.4 Workflow and Schedulers

2.4.1 Kepler

The work of these activities has been focused on building a set of different use cases that can be used as templates for the initial fusion community as well as for 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 distributions of these middleware. A number of different workflows for different scenarios has been implemented and provided to users. Such workflows have been customised for different application use cases mentioned and reported in the activity reports and other deliverables.

As it has been stated in previous reports, the activity has been focused on establishing collaboration with other user groups which have shown interest in our work. The initial group of users and use cases come from the fusion community, which is the main driving force of this activity.

The dissemination activities carried out in different locations and meetings show the impact of the work to other communities. Among dissemination activities we have performed the tutorials for users, made several presentations on community related conferences. As mentioned in the previous report we have established contacts with the computational chemistry community, preparing some of the workflows use cases. The other community group that we are working with is the astrophysics community, which has been provided with the first templates for their use cases, and next use cases are under development.

As concerns sustainability plans, besides the future usage of the implemented use cases and Serpens suite for Kepler inside the supported user communities itself, some of the template use cases implemented in EGI-InSPIRE (as well as new use cases) will be customised and

used in the national PL-GRID+ project for new communities. Initially it concerns the astro community as well as nanotechnology community with number of different scenarios.

2.4.2 SOMA2

SOMA2 is a versatile modelling environment for computational drug discovery and molecular modelling. 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 modelling 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 PQ8 the SOMA2 EGI pilot service was set up and introduced for users in the EGI Community Forum at the end of March 2012. This service enables users to perform virtual screening jobs by using the AutoDock 4 program. This service facilitates resources hosted by the Finnish Grid Infrastructure (FGI). From a development point of view, the main outcome of PQ8 and PQ9 was a new version release of SOMA2. On 31.5.2012 SOMA2 version 1.4.1 Aluminium was released and made available on the SOMA2 web site. In addition to bug fixes, this version adds support for using user generated proxy certificates in grid use via SOMA2. This release also contains grid enabled versions of SOMA2 demo program descriptions, which make use of the Open Babel program package. This version along with the program descriptions were also taken in use in the SOMA2 EGI pilot service. The latter part of PQ9 and start of PQ10 was used for developing general improvements in SOMA2. In addition, CSC has maintained and operated existing SOMA2 service installations.

2.4.2.1 Outlook for future work

During PQ10 CSC aims to publish yet another release of SOMA2 (1.5.0 Silicon) which contains the development efforts of PQ9 and PQ10. It was possible for CSC to employ extra effort in this task, so the project has developed further at a faster pace, resulting in the total use of the allocated EGI SA3 funding. This work is therefore now unfunded and the development effort is focused primarily at the national level. CSC will however support the existing SOMA2 services and it is foreseen that this will also suffice the needs of the international SOMA2 service. CSC can offer some small level of support for SOMA2 communities, but as grid enabled SOMA2 services are advertised to the various user communities, participation and involvement of those communities is encouraged.

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.

As CSC develops and maintains the SOMA2 service nationally 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. As the goal is that some EGI VO's will also take part in supporting the service in the future.

The source code for SOMA2 is open source, so prospective users can easily make a SOMA2 service for themselves locally.

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.

2.5 MPI

2.5.1 Overview

The MPI sub-task focuses 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, are 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.

The EGI Wiki provides the definitive source of information for MPI user support. Improvements to the MPI documentation[4] were made in PY2 and will continue in PY3 as new releases of MPI-Start are developed. The newer release of the UMD glite-WMS and the glite-MPI products contain significant middleware changes that are required to ensure the correct support for many MPI job types. For example, these include:

- User defined allocation of processes/nodes
- Direct Support for OpenMP codes.

Furthermore, end-user support for both MVAPICH (an MPI implementation with advanced Infiniband networking) and support for generic parallel job support is in development and will be tested by SA3 members in the latter stages of PY3. This work is under the leadership of

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

CSIC, and is aimed at implementing part of the "MPI and Parallel Computing" user and technical requirements.

The funded partners of the MPI and Parallel Computing sub-task (CSIC, TCD) do not, in general, develop MPI software for this task. The effort is dedicated towards the provision of MPI support for the large number of user communities that run MPI applications on the Grid. The SA3 team has contributed greatly to the cross-activity MPI Virtual Team (MPI-VT). This team was established in mid-PY2 and ended at the start of PY3. The SA3 team led many of the sub-tasks, and is currently involved in completing the remaining findings/recommendations. Some are expected to be complete after PM29.

The partner UNIPG, who are not funded in PY3, will continue to develop a range of MPI code that has reuse in many other scientific domains including, among others, the Earth Science and Astronomy & Astrophysics communities. UNIPG has also implemented numerous advanced parallel codes for General Purpose Graphics Processing Units (GPGPUs).

2.5.1.1 MPI-Start support for non-MPI parallel applications

TCD has also looked at exploiting the MPI-Start framework maintained and developed by CSIC under EMI to run non-MPI parallel codes at resource centres. For example, they have shown that it is possible to run Charm++ codes using MPI-Start as the mechanism for launching distributed processes. This has very good potential in supporting non-MPI parallel applications, such as some Map-Reduce like models, and will be investigated further in PM29 to PM36.

2.5.1.2 General Purpose Graphics Processing Units

Over the past year, there has been a substantial increase in user communities wanting to exploit the parallel processing capabilities of General Purpose Graphics Processing Units (GPGPUs). At a community level, UNIPG has been actively developing its CHIMERE code to run on GPGPUs. Both UNIPG and TCD are working on exploiting GPGPUs. The UNIPG work is already in production: 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. This work will continue during PM29 and PM36.

TCD also leads a new EGI Virtual Team – the GPGPU-VT. This VT has surveyed User Communities to gauge current and planned GPGPU usage in the Grid environments. The GPGPU-VT reported its findings at the EGI Technical Forum 2012 in Prague (PM29).

Of the 44 resource centres which responded, 13 (30%) currently provide GPGPU resources, and of these 85% intend to increase capacity. 23 (>50%) of the surveyed resource centres intend to provide GPGPU resources in the next two years. Most (will) provide access via Torque/Maui, and fewer use Slurm or another LRMS. Nvidia dominates the GPGPU systems in use.

In the user survey, there were 47 respondents, 39 of whom expressed interest in remote GPGPU access via EGI or NGIs. Just over half of respondents currently use GPGPUs. Users

would like exclusive access to GPGPU WNs; and CUDA is the preferred development platform.
We expect that there will be some user community follow-up work during PM29 and PM36.

2.5.2 Sustainability

The sustainability of a MPI-related activity beyond EGI-InSPIRE can be broken down in two distinct key parts:

- sustainability of MPI grid support teams, and
- sustainability of MPI-related software components.

MPI user-community support is currently handled as a dedicated task under SA3.
This task is due to end in PM 36. The task has also received continued unfunded support from the University of Perugia (UNIPG), whose members have freely contributed to the MPI task beyond their original contractual period. The MPI applications developed and implemented by UNIPG, namely reactive scattering and secondary pollutants production codes, shall be maintained to keep running on the grid. In addition, the short-lived MPI Virtual Team provided a significant boost in delivering much needed effort over the project year 2. This has resulted in improvements to user and resource centre MPI documentation, as well as significant performance and reliability improvements at the resource centres. The MPI Virtual Team was composed of members from the EGI-InSPIRE SA1, SA2 and SA3 activities, and has shown how continuous expert input is required from a varied range of users, resource centres and software providers.
The sustainability of the SA3 MPI activity beyond EGI-InSPIRE is unclear. Some of the participating institutes are under severe financial pressure, and this may persist beyond the lifetime of the EGI-InSPIRE project.
A similar future activity would probably have to be redefined as a smaller core coordinating task that depends on committed user communities, NGIs and resource centres to provide the required support structures. However, for MPI activity to exist and grow within EGI, this task should be retained and handled by one or more of the expert project partners.
The EGI MPI user community depends on several external software providers. These include the MPI implementation providers, such as MPICH and OpenMPI, as well as MPI-Start, an intelligent software suite that is developing a much needed uniform access to multiple MPI and grid middleware implementations.
The sustainable provision of these MPI software components can be considered somewhat positively:

- Both MPICH and OpenMPI have good commercial support from several large academic and research institutes. Further financial and technical support is received from hardware vendors.
- MPI-Start it currently developed within the FP7 EMI project. It is evolving to work with the many different grid middleware implementations supported by EMI and deployed on the EGI infrastructure.

Current development includes extending support to the MVAPICH2 MPI suite.
This is widely used fault-tolerant implementation of the MPI reference API. Additional developments include improved support for the batch scheduler system SLURM. Post EMI, the maintainers of MPI-Start (CSIC) hopefully expect to continue development within the the context of the ScienceSoft initiative. However, the future status of the ScienceSoft initiative is

not yet known. To assuage such problems, and should ScienceSoft not come into fruition, the MPI-Start software packages will also be available through the widely used EPEL Software repository and distribution channels. This also has the added benefit of making it available to a much wider (non-grid) audience.

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 CRAB and CRAB Analysis Server

3.1.1.1 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 been in production and in routine usage since 2004 and it has demonstrated being an important production service to publish physics results. During 2011 CMS observed about 250 unique users a day, submitting about 200K CRAB jobs per day, with close to 1000 individuals per month. The main development items achieved between the end of 2011 and first half of 2012 include: the re-implementation of a new distributed monitoring data service based on NoSQL databases that guarantees to collect and preserve all needed information; the introduction of an automatic user data publication system to optimize user interaction with the data management services; the development of a new RESTFul layer between the CRAB client and the server which improves the interaction with the analysis system, allowing introduction of an authentication and an authorization layer, providing the functionalities needed to aggregate all the distributed information through a well defined set of user resource oriented APIs. Initial scale tests did not show any major issues in the CRAB dedicated components. A second phase of testing has also included the CMS integration team plus beta-users to provide feedback on the system and in particular on the information provided through the new monitoring and APIs.

Main working items for 2012 and the beginning of 2013 will include:

- Continuing to perform scale testing, in particular focusing on the last two steps of the analysis workflow; the asynchronous copy of user analysis results and the automatic publication system. This phase will provide an understanding of the scalability of the system before really exposing the service to the final users.
- Feedback-driven development in order to converge as closely as possible to user requirements, optimizing in particular the user point of view of tracking their requests. This step will be an iterative process between developers, the integration team and beta-users.
- Starting to move the CRAB dedicated components in the CMS production environment in order to facilitate the testing and to match the Service Level Agreements defined by the experiment.

3.1.1.2 Sustainability

CRAB is a joint project involving 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 60th 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. Furthermore, the main CRAB components have been included in the feasibility study and proof of concept phases of a wider project which aims to evaluate the possibility to provide a common analysis workload management system among different LHC experiments.

3.1.2 Data Management

3.1.2.1 ATLAS Distributed Data Management

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 ATLAS DDM software developed through EGI InSPIRE SA3.3 funded effort is in a mature state and was handed over to the ATLAS DDM team from the CERN-PH-ADP group, whose present work is limited on maintenance and support operations of this software.

At the same time, the CERN-PH-ADP group is designing and developing a new DDM system (the Rucio project [RUCIO][5]) to solve current shortcomings and scalability issues in the current architecture, as well as to evaluate the integration with new technologies (e.g. object storages).

3.1.2.2 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. 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 addition, EGI InSPIRE TSA3.3 funded manpower implemented a system for the dynamic cleaning of unused replicas when the different physics groups exceed their allocations on sites. Both systems are in production and will be presented in the CMS Week at the beginning of September. The outcome of this meeting will define if there is future work needed on these systems.

[5] Rucio: Conceptual Model; M. Barisits et al, ATLAS internal note reference: ATL-COM-SOFT-2011-030, http://cdsweb.cern.ch/record/1386633

3.1.3 Persistency Frame-work

3.1.3.1 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 communities other 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.

3.1.3.2 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. 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. For 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. 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.

For 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 et al.) 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).

The Frontier/Squid system is also relevant to CORAL. For CMS, the Frontier/Squid deployment is monitored and supported by a small team composed primarily of institute-contributed effort. The system is also now becoming more critical 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.

3.1.4 DIRAC

3.1.4.1 Overview

The DIRAC system provides a complete solution for using the distributed computing resources of the LHCb experiment. DIRAC is a complete framework for data production and analysis, including workload management, data management, monitoring and accounting (more details have been given in document [MS610]). LHCb DIRAC is the DIRAC extension for the LHCb community, which has been formally separated from the more general framework in order to include developments specific for the LHCb VO. One of its most important components is the Data Management System (DMS) which the EGI-InSPIRE project has supported since October 2010. During the last semester considerable effort has also been put into improving the Production (or Transformation) System (TS).

3.1.4.2 Sustainability

The roadmap for the next months includes the following items:

- Consolidate the popularity service that was developed during the first half of the year and put into production during the last quarter. Such a service should provide metrics to assess data popularity and provide a ranking of the most popular datasets (i.e. those most frequently accessed by users). The final goal is to use data popularity information to implement a dynamic data placement model, where the number of replicas of a given dataset is related to its popularity. This would considerably help the VO's data managers to optimise the usage of storage resources on the grid and would automate many operations that are currently done manually.
- Support and possibly improve the system for the accounting of storage resource usage that was developed and put into production during the past year. This system, already extensively used by data managers, has recently undergone some improvement, taking advantage of new features of the LHCb DIRAC DMS, and a new version is currently being validated. It provides powerful and flexible accounting for storage resources, which is crucial for planning data management operations and optimising the usage of storage at grid sites.
- Provide general support for LHCb computing operations on the grid, both for production and private user activity. In particular, since the restart of data taking in April a considerable effort has been dedicated to production management. A more sustainable approach to production management is being promoted, including efforts to define good practices, enhance monitoring and define common procedures. This will aim to make the data production operations more sustainable and less expensive in terms of man power.

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. The LSGC implements the concept of VRC in EGI. VRC management is a complex duty for which appropriate support tools need to be developed: integrated tools are needed to provide a VRC-wise vision of the activity and facilitate the VRC administration, help VOs of the community to mutualise efforts and leverage common tools to avoid duplicating efforts. As described in MS617 [MS617], the Life Sciences HUC contributes to the LSGC effort to organize the community and deliver those services.

3.2.2 Delivery status

In this context, several services to the HUC users have been designed and provisioned. Details can be found on the LSGC wiki[6]:

Communication services:

[6] LSGC wiki: http://lsgc.org/en/LSGC:home

- Per-VO and VRC-wide mailing lists have been set up and are kept up to date to ensure communication within the community.
- Several web gadgets customized for the Life Sciences have been added to the LSGC wiki, with the help of the User Communities Support Team: Life Sciences applications from the Applications Database[7], community requirements posted to the Requirement Tracker[8] system, Life Sciences trainings from the Training marketplace[9]

Technical support services:

A HUC technical support service, delivered by a team of expert users, has been set up to address the difficulties reported by users or resource providers (sites). Phone conferences help teams relay each other during duty shifts. See technical team wiki for details[10].

The technical support team has deployed tools to monitor the infrastructure health at a VO level:

- A dedicated Nagios server[11] deployed by the French NGI;
- Community files management: monitor storage space consumed VRC-wise, anticipate problems of storage resources filling up, handle SEs decommissioning, file migration procedures.
- Centralized view of VO resources that are currently not up and running (downtimes, not in production...);
- Miscellaneous tools for facilitating daily follow-up of issues, manual checks etc.

During the last year, the LSGC 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 revised as detailed in section 2.3.1.2.

Previous plans to deploy a redundant LFC server have been abandoned for several reasons detailed in [MS617]. To summarize, the effort required to build a redundant LFC configuration would not be worth the interest it would bring.

3.2.3 Sustainability and future plans

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

As explained above, a VRC-wise vision of the activity is needed to help the community to mutualise efforts and leverage common tools, thus avoid duplicating efforts. Among existing projects, the VO Operations Dashboard[12] (part of the EGI Operations Portal[13]) is a promising

[7] http://appdb.egi.eu/
[8] https://rt.egi.eu/rt/
[9] http://www.egi.eu/services/support/training_marketplace/
[10] http://lsgc.org/en/Biomed-Shifts:Index
[11] https://grid04.lal.in2p3.fr/nagios/
[12] https://operations-portal.egi.eu/voDashboard
[13] https://operations-portal.egi.eu

tool expected to become the daily tool for VO operation support teams. It covers important operational needs from the LSGC point of view such as visualization and classification of alarms from monitoring sources (Nagios), integration with the GGUS ticketing system to assist in the creation and follow-up of team tickets, cross information with resource status from GOCDB, support shifts organization and takeover reports, as well as possible advanced features like the detection of flapping service or report builder.

Taking the biomed VO as a significant use case, the Life Sciences Grid Community (LSGC) directly interacts with the development team of the VO Operations Dashboard, suggests and discusses features, and provides feed-back as a beta-tester.

In addition, the LSGC envisages the development of an LSGC dashboard[14], which has been delayed due to the end of the HealthGrid association. The dashboard is expected to integrate various VO management services into a single portal:

1. User management tools, covering users life cycle management, interface with the VOMS and Application Database services;
2. VRC-wide accounting, needed to deliver statistics at the VRC level. The current EGI accounting portal only provides per-VO accounting information
3. Existing tools listed in Technical support services in the previous section;
4. Monitoring of computing resources availability to detect possible bottle necks among computing resources, and eventually address this issue at different levels, for instance in job submission policies.

Last, the LSGC holds an ongoing discussion with storage resource providers and the HEP community to identify best practices in terms of community data management, and extend procedures and tools developed so far to more general procedures tailored for small communities like the LSGC VOs.

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 March 2012 to September 2012 and the plans for the period from October 2012 to April 2013 have been focused on the following topics:

- Coordination of activities related to the usage of distributed e-Infrastructures by the A&A community;
- Tools and Services for data visualization: VisIVO;
- HPC, parallel programming and GPU computing;
- Access to databases from DCIs and interoperability with the VObs (Virtual Observatory) data infrastructure;
- Long-term sustainability plan.

[14] http://lsgc.org/en/LSGC:lsgc-dashboard

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 like SKA (Square Kilometer Array) and CTA (Cherenkov Telescope Array) continued also in the period covered by ms620.

A special effort is now in progress to gather as many astronomical requirements, applications and workflows as possible given that this is the main contribution requested to the community by several projects where astronomers have a direct participation and whose goal is to port workflows coming from different communities on specific platforms. This effort is in line with the outcome of the astronomical workshop organized in Paris in November 2011 and with what emerged during the EGI User Forum 2012 in Munich.

To achieve this goal contacts have been already established with groups and Institutes recognized as potential contributors of workflows and applications. An informal meeting has been organized in the framework of the EGI Technical Forum 2012 to meet interested people and go ahead with the applications/workflows gathering process.

According to the outcome of the workshop held in Paris, the identification of new astronomical VOs (Virtual Organizations) and VRCs to create in EGI will also continue through direct contacts with interested groups/Institutes.

The activation of a new community group in OGF, called Astro-CG, was also submitted and approved during the Summer 2012. Astro-CG will act as the place where astronomers distributed worldwide meet, discuss and propose solutions and standards for what concerns distributed infrastructures (both data and computational) and ways to interface them with a special attention deserved to the Virtual Observatory for what concerns data e-Infrastructures. Very close contacts will be established between the Astro-CG in OGF and EGI.eu.

3.3.2.2 VisIVO

After the production and the deployment of the first grid-enabled version of VisIVO, based on a specific grid-enabled library that allows users to interact with Grid computing and storage resources and able to interface with and use the gLite Grid Catalogue, further developments concerning VisIVO are in progress; they are those planned in writing the ms614 milestone report, namely: 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.

As already highlighted, VisIVO was initially conceived and implemented as a visualization tool for astronomy, but recently it evolved in a generic multi-disciplinary service that can be used by any other community that needs 2D and 3D data visualization.

3.3.2.3 Grid and HPC

As stated for milestone ms614, activities related to Grid and HPC are carried out in close coordination with EGI and with IGI (the Italian NGI) in Italy; such activities continued also during the period covered by ms620. We remind here that one of the goals of this activity is the introduction of small-size HPC resources in Grid and then the installation/configuration of

HPC clusters based on low latency/high throughput networks, HPC libraries and tools, modules and compilers on one hand and on the other hand make the grid middleware aware of these resources. A&A applications able to take advantage of such small-size HPC resources have been already identified and now they are going to be used as testers and validators of these resources.

GPUs as well are also emerging as important computing resources in Astronomy and the option of using those offered by Grid sites to make visualization processing on VisIVO is currently under evaluation. As planned at milestone ms614, this activity is going ahead in close coordination with all other communities of IGI in Italy. During the last period of SA3, GPU-related activities undertaken by EGI.eu will be also monitored looking for common plans and possible joint activities.

3.3.2.4 Access to Databases and interoperability with the VObs

The identification of use-cases and test-beds (both applications and complex workflows) that require simultaneous access to astronomical data and computing resources directly impacts the topic related to the interoperability between data and computing e-Infrastructures and then the interoperability with the Virtual Observatory. The activity in progress concerns first of all the implementation of a SSO (Single Sign On) mechanism able to grant access to computing and data resources through a single authentication phase and then tools/services to access astronomical databases federated in the VObs from DCIs. The deployment of SSO solutions and of tools/services enabling the access to databases, in turn, are closely related with the implementation of smart portals and science gateways using these tools and services. The creation of a VT (Virtual Team) to design and implement a science gateway and the related SSO service for the CTA project is currently in progress.

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 most recent evolution, however, it is now a general purpose visualization tool, able to meet the requirements of many communities and of their applications. A sustainability plan for VisIVO has recently developed, based on a community of end users that is continuously expanding and on the fact that the core of the development/maintenance group is formed by staff personnel of INAF, the Italian Institute for Astrophysics, who will continue to invest on VisIVO in the next years. The sustainability plan set up for VisIVO has been presented at the EGI Technical Forum in Prague.

As highlighted in the section dedicated to the coordination of the A&A community, the most relevant activity currently in progress is the identification of A&A projects and Institutions able to bring contributions in terms of astronomical requirements, applications and workflows and available to create and coordinate new VOs and VRCs focused on specific topics; in this perspective the preparation and implementation of a sustainability plan should become easier.

3.3.4 Future Plan

The plan of activities covering the next period before the end of SA3 (April 2013) aim at:

- completing the activities currently in progress for VisIVO to produce a version of the tool fully compliant with DCIs (in particular with the European DCI based on EMI middleware)
- progressing with the requirements, applications and workflows gathering process to collect as many material as possible to strengthen the strategic weight of the community in the context of distributed e-Infrastructures and maximize the contribution of the community to key international astronomical projects.
- contributing to projects and initiatives for a combined use of computing resources of different nature and a full interoperability between them and the data e-Infrastructures.
- designing a long-term sustainability plan for astronomical tools and services developed in the context of international projects and initiatives in a way that their long-term maintenance is guaranteed and people can adopt them safety for their scientific goals.

3.4 Services for the Earth Science Virtual Research Community

3.4.1 Progress and plans

The Earth Science VRC develops and maintains tools and interfaces to data outside of EGI, more specifically data from Earth Science specific data sources. The applications in the community typically need original data from data centres, which is delivered through storage servers using specific protocols or legacy protocols such as HTTP, in different formats, which often makes discovery and access from the Grid cumbersome. The GENESI-DR Infrastructure from the GENESI-DR and GENESI-DEC[15] FP7 projects led by ESA can be of assistance in this situation. The projects in this task include a command-line client for the above mentioned GENESI-DR infrastructure (gsearch), a web GUI leveraging this client in Grid jobs, a command-line tool for the comfortable download of ESGF data (synchro-data) and a Credential Translation Service (CTS) for binding EGI grid credentials to ESGF credentials. Regular adaptions are necessary due to changes at the level of the GENESI-DR protocol returned by the search services and its connected data repositories. As the GENESI-DEC project has finished now, we expect more stable results in the future. The VRC extends the discovery and access possibilities to other providers and different protocols by using an instance of GI-cat, a web service that can broker and federate external catalogues, developed by ESSI-Lab (http://essi-lab.eu/gi-cat).

The following tasks have been identified for the command-line client "*gsearch*" and are reasonable to improve both quality and service:

- Further adaption to the GI-cat web service and the possibility to download data from a variety of data centres directly to a computing or storage element
- A timeout and watchdog handler for controlling and managing calls to different sites, which are not available or only reply with high latency. An improved threading system will ensure that the data from servers with high latency will be processed without impact on the Text User Interface (TUI), or the other search instances.

[15] http://www.genesi-dec.eu/

- Enhanced handling of a wide variety of optional XML OpenSearch attributes, for all kind of different data. For instance, satellite data with an orbit number tag is useful, but not for data from a research ship on the Mediterranean Sea.
- Development of an independent light-weight download tool without dependencies (meaning necessary dependencies are statically linked) with support of common protocols like http, https, OPeNDAP…

Initiated by requirements from ES users, a portal solution with a flexible web GUI was developed and is further maintained. For this part, the following tasks have been identified:

- **Full portal integration**
  The web GUI is currently available as a portlet. At the moment, the integration activities (as mentioned in QR 9) of the portlet targeting existing Grid portal framework (such as WS-PGRADE[16] and the Vine toolkit[17]) are on-going. Besides exploiting existing certificate and user management techniques, further work is oriented at job monitoring and storage handling.
- **Modifications due to new GI-cat releases and changes in the attached catalogues**
  The current GI-Cat version, which is used as a backend catalogue provider for the gsearch web interface, is 9.3-beta-1. So far, gsearch has not been adapted to work with newly introduced OpenSearch parameters. Another main issue and by thus a main work, are changes of the attached catalogues. We assume that this work will be going on until the end of the project.
- **jLite Job Submission and JDL**
  The gsearch portlet uses the jLite Java-API for the JDL submission. Unfortunately, jLite does not seem to be maintained anymore. One part of the work will thus be the identification of possible problems because of updates. The second facet is the job itself and customizations by users. A more flexible job creation mechanism will be designed that is easily customizable for users and exploits the download tools developed in the task to transfer the data sets returned by the prior search.

The intelligent ESGF data transfer tool, named "*Synchro-data",* that facilitates the command line, bulk oriented access to ESGF data was updated. The tool can download files from the ESGF infrastructure in an easy way, through a list of variables, experiments and ensemble members. Planned developments include:

- Enhanced documentation with usage examples
- Implementation of GridFTP downloads from ESGF data nodes
- Improved multi-project support (originally the tool was CMIP5 specific)

[16] http://www.guse.hu/
[17] http://vinetoolkit.org/

- General ergonomic enhancements
- Increased stability and reliability
- Add option to remove 'waiting' transfers
- Retrieve models list from search-API

The Credential Translation Service needs additional documentation and broader dissemination. It will also be updated to recent MyProxy versions.

3.4.2 Sustainability and cross-community aspects

As the developed software is rather focused on its functionality - it consists of atomic applications developed for specific use-cases of the Earth Science community - we fear that the prospective gains of reuse for other communities are rather limited. Nevertheless, the inquiries and technical developments may be of assistance, or even serve as groundwork for other communities and their respective application areas. Despite their varying technical conditions (e.g. concerning data handling tools), 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. The “CTS” can be used by other communities as is. It is freely available as a patch to the MyProxy credential management system and can be used as a transitional solution until an official EGI solution is available, such as the planned EMI Security Token Service.

4 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
MS617	MS617 - Services for the Life Science Community https://documents.egi.eu/secure/ShowDocument?docid=1289
Shamir	A. Shamir. "How to share a secret" in Communications of the ACM (CACM), vol. 22, pages 612-613, 1979.