EGI-InSPIRE

HUC SOFTWARE ROADMAP

EU MILESTONE: MS606

Document identifier:	EGI-MS606-V5
Date:	**16/05/2011**
Activity:	**SA3**
Lead Partner:	**INFN**
Document Status:	**FINAL**
Dissemination Level:	**PUBLIC**
Document Link:	https://documents.egi.eu/document/310

Abstract

This document provides the second 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.

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 are 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.



I. COPYRIGHT NOTICE

II. DELIVERY SLIP

	Name	Partner/Activity	Date
From	Laura Perini, Jamie Shiers	INFN/CERN	14/02/2011
Reviewed by	**Moderator:** Pekka Lehtovuori **Reviewers:** Onurt Temizsoylu	CSC TUBITAK ULAKBIM	14/03/2011
Approved by	**AMB & PMB**		15/05/2011

III. DOCUMENT LOG

Issue	Date	Comment	Author/Partner
1	01/02/2011	First draft	Laura Perini, Jamie Shiers
2	15/03/2011	Second draft	Laura Perini, Jamie Shiers
3	15/03/2011	Final	Laura Perini, Jamie Shiers

IV. APPLICATION AREA

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

V. DOCUMENT AMENDMENT PROCEDURE

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

VI. TERMINOLOGY

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

VII. PROJECT SUMMARY

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

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

EGI-InSPIRE will collect user requirements and provide support for the current and potential new user communities, for example within the 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 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 Dashboards, GANGA and related applications, the HYDRA and GrelC services, the Kepler, Gridway, SOMA2 workflow schedulers, and the enabling of Message Passing Interface (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 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 the European Middleware Initiative - EMI).
- To use the experiences obtained by these early adopting communities in integrating new data sources, tools and services 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 from their current e-Infrastructure provider.

The report and planning of the services and tools, sketched in this report, shows relevant developments and widespread usage, mainly still concentrated in one or two communities. The relevant effort going toward making easier the use of these services and tools, 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 it.

TABLE OF CONTENTS

1 INTRODUCTION

This document provides the second 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 (HUC) 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 can be interested in adopting parts of the software for their uses and in suggesting developments to this effect.

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 toward a single community, however much effort is devoted to the documentation and to the easiness of use of the different products, both 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.

2 ROADMAP FOR THE SA3 SHARED SOFTWARE SERVICES AND TOOLS

2.1 Dashboards

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. The system is not coupled with the particular data management or workload management systems used by the virtual organizations (VOs). Instead it offers generic monitoring solutions working across several middleware platforms. The Experiment Dashboard is widely used by all LHC VOs, in particular by ATLAS and CMS. The LHC VOs are involved in the development and validation of the Dashboard applications. They also take part in operations of some of the Dashboard servers, for example the ATLAS DDM Dashboard. Most of the Dashboard applications are generic and are shared by several LHC VOs.

The development of the Experiment Dashboard system is mainly driven by the requirements of the LHC user community. However, many of the Dashboard applications are generic and can be used outside the scope of LHC and HEP. For example, Site Usability and Site Status Board may be of interest to any Heavy User Community which uses a lot of distributed resources and needs to monitor their state and usage. All Dashboard applications are written within a common framework, which provides the building blocks for development of the monitoring tools. The framework is well documented and the developer's guide includes detailed descriptions of the framework components and usage examples. Therefore adding new functionality to the existing applications or adapting them for other VOs is straightforward.

The main development effort in 2011 will focus on a partial redesign of some of the Dashboard applications leveraging new technologies such as Asynchronous JavaScript and XML (AJAX). The aim is to provide faster, richer and more intuitive user interfaces. The migration of the cluster of the Dashboard servers to the virtual machines started in 2010 and will continue in 2011. This will allow to decrease maintenance and support effort and to improve efficiency of the Dashboard deployment model and reliability of the Dashboard servers. Another goal consists of extending the functionality of the Dashboard applications following the requests of the LHC user community, as described below.

Main tasks in 2011:

- Development and deployment of the Site Usability portal that will be compatible with the new architecture of the Service Availability Monitor (SAM). SAM was recently redesigned and the new version is based on the Nagios framework.
- Development and deployment of the new version of the Distributed Data Management Dashboard for the ATLAS experiment, which will make it easier to identify the location of failures in the data transfer chain. The key feature is an interactive matrix of transfer statistics showing both source and destination with links to error samples and transfer details.
- Extending the functionality of the Dashboard job monitoring in order to facilitate daily operational tasks, among them user support for analysis users, efficient data replication, and

cleaning of the disk space at the distributed sites.

2.2 Applications

2.2.1 GANGA

Ganga [GANGA] is an easy-to-use frontend for job definition and management, implemented in Python. Ganga is a mature product with 500 unique end-users every month. It is the main user tool for running grid analysis jobs in LHCb and ATLAS experiments. CERN leads the Ganga Core development and contributes to ATLAS-specific Ganga functionalities.

The Core part of Ganga is a software code-base which is shared between different user communities. The Ganga Core software is driven by the evolution of requirements of Heavy User Communities (HUCs), and High Energy Physics (HEP) HUC in particular. This is performed in several ways:

1. Missing capabilities and features of general interest are developed in the Core software and are made available to HUCs,
2. Core software is adapted to the needs of HUCs by providing better foundation and integration for HUC-specific features,
3. HUC-specific features that are of general interest are refactored into the Core software to increase code sharing and consistency of the Core software environment

Specific work items in Ganga Core include:

- Refactoring of application configured state to capture existing functionality in the 'GangaAtlas' plugins and provide a common way to support it
- Further consolidation of task monitoring capabilities based on Task Monitoring Dashboard for Atlas HUC
- Consolidation of project web pages and update of project documentation
- Integration of the software into standard Unix distributions

Additionally, the ATLAS experiment modules will undergo some evolution in the next year:

- The Ganga-ATLAS module will be factorized to reflect the support for the PanDA backend as the main backend for ATLAS. This includes the features currently available in other backend plugins used by ATLAS to be ported to the Ganga-PanDA backend plugin.
- Improved task bookkeeping system will be developed to more intelligently present the overall status of analysis tasks to the users (PanDA jobs are often retried and rebrokered to other sites which makes the bookkeeping of an analysis task rather tedious).

In addition to the above, there are significant contributions to Ganga Core that are not funded by EGI-InSPIRE and which include improved support for automatic merging and resubmission and refactoring of ATLAS tasks into common framework.

2.3 Services

2.3.1 HYDRA

Hydra is a file encryption/decryption tool developed as part of the gLite middleware. The plans to deploy a Hydra catalogue for the Life Sciences (LS) user community starting from year 2011 are unchanged. The scheduled roadmap for this service provision is:

- January 2011 (M9): installation of a test Hydra catalogue and validation of the functionality delivered, including:
    - Encryption key registration and removal.

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

- April 2011 (M12): delivery of a distributed Hydra catalogue for production use within the LS area.
- Unless specific intervention is required due to changes in the gLite middleware Data Management System, a functionality check is then planned every semester from month 15:
  - July 2011 (M15): in preparation for D6.4, revalidation of the functionality delivered with performance and scale tests (encryption overhead measurements, number of keys that can be stored and number of accesses that the Hydra catalogue can handle).
  - January 2011 (M21): revalidation of the functionality delivered.
  - July 2012 (M27): in preparation for D6.7, revalidation of the functionality delivered with performance and scale tests.
  - January 2013 (M33): revalidation of the functionality delivered.

2.3.2 GRelC

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.

The aim of this data grid service is to efficiently, securely and transparently manage databases on the grid, across virtual organizations. The GRelC service provides dynamic binding to different versions of MySQL, PostgreSQL and SQLite for the relational data model as well as to eXists and Xindice for the hierarchical one. Support is also available to Java and C developers (in terms of JDK/SDK, documentation and examples) through the GRelC web site (http://www.grelc.unile.it).

During the project, the GRelC system (the network of GRelC services deployed within EGI) will be enhanced to support the EGI communities with a new set of functionalities available as a web application through the GRelC Portal and the DashboardDB (old name Dash-G) interface.

In particular the new releases of the GRelC Portal will help users to easily and transparently port in grid existing DBs, manage several GRelC service instances, grid users, data access policies, etc.

Support in terms of management, monitoring and control of the GRelC services will be provided through the DashboardDB (old name Dash-G). The DashboardDB will provide several views (including charts, reports, tables) about the GRelC deployment, the status of the services, the list of available grid-databases, the supported VOs, related keywords/tags, user's comments, etc. The design of this web application started at the end of Q2 and continued in Q3. Some internal modules (Java classes) were also developed in Q3. The implementation will continue in Q4. A preliminary version (v0.1) with the home page will be available at the end of Q4 in order to highlight the main goal of the DashboardDB and start collecting feedback.

The infrastructure behind the DashboardDB has been designed, implemented and tested during Q1 and Q2 and it consists of:

1) a set of sensors developed in C to collect important metrics like the grid-db service availability, the service failure and failure type (service down, service unreachable, etc.) and round trip time from a central location (the DashboardDB machine), other statistics and data summaries;
2) a relational database (named *system catalog*) already in place, where all of the metrics collected by the sensors developed in 1) are stored.

The future plan related to this part of the DashboardDB infrastructure basically focuses on doing tests and bug fixing. No major changes are foreseen.

Going to the web application part, during the next 6 months the DashboardDB will be available

online with the following features: **project** and **host** views. This means that a set of internal modules for user management, security, project and host management will be developed as well in order to support the two main views (v1.0).

The work on the **service** view and on the **grid-database registry** (for the EGI Database of Databases) will start during the second year. By the end of the second year a second major release (v2.0) of the DashboardDB views (project, host and service) as well as the grid-database registry will be available online. The registry will contain all of the information about the grid-databases available in the GRelC system, the associated VOs, a description, some keywords and other useful metadata. During the third year of the project, the final release (2.2) will basically offer little new functionality, but a higher level of robustness. Comments, feedback and requirements coming from the user communities will be taken into account (and will be essential) to improve the software from a functional and non-functional point of view.

The registry will complement the functionalities provided by the EGI Application Database and will represent a distributed and multi-VO system. In few words, the DashboardDB will provide a *cross-VO grid-database system* giving the user community both a unified view about the available grid-databases and several tools to interact with them.

The design of the internal modules of the DashboardDB web application performed during Q3 took into account the **Web2.0** paradigm. “Mashup”, Google Maps, permalinks, comments, are just some of the features that were considered during the design phase. During the next 6 months the implementation will proceed accordingly, implementing “export” (permalink) capabilities and “map-based” (through the adoption of Google Maps) interfaces for projects and hosts views.

An important task started during Q2 and completed during Q3 was the SA3 EGI Questionnaire. The questionnaire aims at providing an up to date list about databases (relational, XML-based, etc.):

- Already in place but that need to be ported to the grid in the context of the EGI-Inspire project;
- Already ported to the grid and therefore accessible in the context of the EGI-Inspire project;
- Not yet deployed in the context of the EGI-Inspire project;
- Available from external sources via FTP, HTTP and that would need a grid-enabled instance and interface in the context of the EGI-Inspire project.

The questionnaire was sent to the HUC during Q3 and preliminary feedback is expected starting from Q4. As a follow-up of this census, a preliminary static list about the available database resources in the context of EGI will be inferred starting from Q4. The list will help to define use cases, understand needs and collect requirements from the HUC.

Concerning the Earth Science community, the Climate-G testbed will represent one of the test cases related to the GRelC service. In this testbed, the GRelC services will be exploited to manage metadata information distributed across several countries and stored into both relational and XML databases. New requirements in terms of metadata management coming from this community (in particular the harvesting functionalities) will be collected and implemented starting from Q4. Support will be also provided through the Climate-G portal, the scientific gateway of the testbed that provides search and discovery functionalities, metadata web pages, lists of experiments, datasets, projects, access to data, etc. Additional use cases related to other communities (e.g. Life Sciences) will be defined jointly with them during the project. For the next months, starting from the input collected through the SA3 EGI Questionnaire, the GRelC team will provide support to implement some specific use cases related to the porting in grid of existing LS databases. Updates on the GRelC software will also reflect new requirements coming from the LS community.

2.4 Workflow and Schedulers

2.4.1 Kepler and Gridway

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 also provides the **runtime mechanisms** needed for dynamically adapting the application execution

As indicated in MS602 [MS602], a list of possible workflows among applications is shown in Table 1.

1) VMEC + DKES	2) VMEC + ISDEP	3) FAFNER + ISDEP
4) ASTRA + TRUBA	5) ASTRA + TRUBA + FAFNER	6) EUTERPE + ISDEP
7) GEM + ISDEP	8) ASTRA + GEM	9) ASTRA + GEM + TRUBA

Table 1 List of possible workflows.

All the workflows have been analysed and designed in detail and work needed for running each of them has been evaluated thoroughly. Basic services needed for the execution of such workflows have been deployed and configured.

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

Among these workflows, the Astra – Truba (4) has been built (see Figure 1). This workflow shows the coupling of two codes running on different computational environments (ASTRA running on HPC and TRUBA running on the grid).

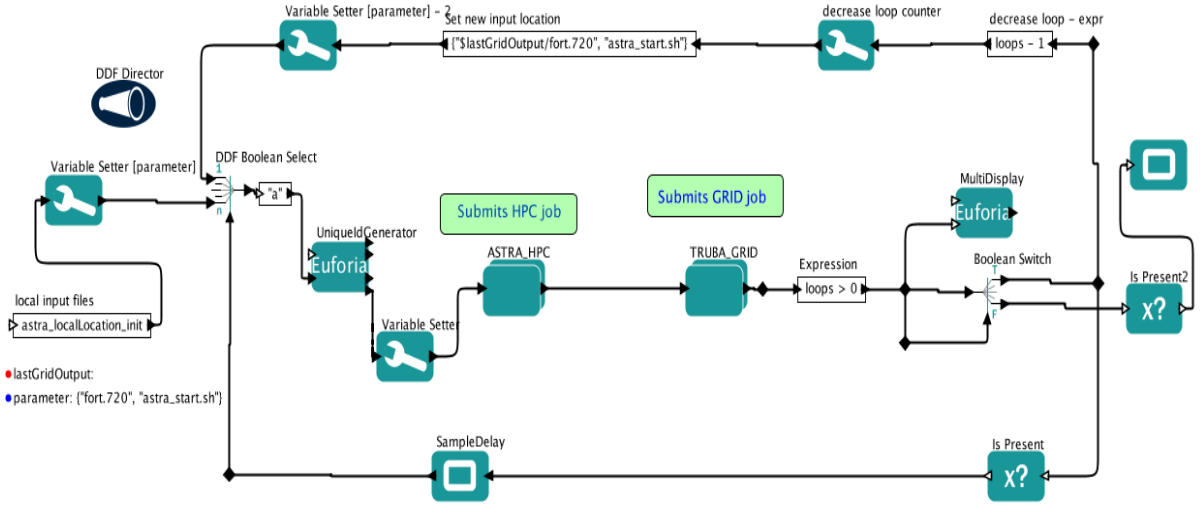

Figure 1 – Astra (HPC) + TRUBA (GRID) workflow built in Kepler.

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

The workflow already built does not include the use of GridWay. However, some of the proposed workflows will require this metascheduler. Thus, the support of GridWay from Kepler must be developed. This support or integration has been already designed and will be implemented in coming months.

As part of the possible exploitation of the currently developed workflow the Computational Chemistry community has been contacted in order to start a possible collaboration in terms of usage of Kepler for some of their applications. As well a Kepler workshop/ hands-on tutorial has been proposed and submitted for the upcoming EGI User Forum. This tutorial includes the basic usage of Kepler and also more advanced and debugging examples will be shown making usage of the grid middleware. All the workflows shown in Table 1 are suitable to provide relevant physics results and their exploitation will be carried out based on this relevance and taking into account the experiments being executed by the involved researchers.

2.4.2 SOMA2

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

During the Q3 of the first project year, CSC has started to investigate suitable application candidates to be integrated with SOMA2. Also we have maintained and operated CSC's SOMA2 service. On the program development, no progress has been made during the Q3.

For the upcoming six months, we plan to release a grid-enabled version of SOMA2. Also, we intend to create a SOMA2 capsule for the chosen pilot application. This will enable to set up a grid-enabled application service, initially on a national level, in CSC's SOMA2 provided for Finnish academic

researchers. This has been our target for the first project year. In Q1 of the second project year we intend to operate and support the SOMA2 service. In addition, we will advertise the upgraded SOMA2 service to existing user communities.

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

2.5 MPI

The MPI sub-task focuses on a number of objectives over a 36 month period. This sub-task is producing numerous MPI workbenches of increasing complexity with specific high impact on the Computational Chemistry, Fusion and Astronomy and Astrophysics (A&A) communities. These products are also intended to have an impact on other user communities. In addition, it focuses on ensuring that the user communities and site administrators benefit from several rudimentary improvements to the methodologies used and the available documentation. Many of these objectives are iterative, often requiring updates or fine-tuning. Other objectives, such as the EGI User Forum participation and the EGI technical Forum, will be repeated at regular intervals. The core sub-task objectives (which bring definition to the tasks sustainability) are:

- Improved end-user documentation, addressing MPI application development and job submission in ARC, gLite and UNICORE
- Quality controlled MPI site deployment documentation
- Outreach and dissemination at major EGI events and workshops
- User community, NGI and site engagement, gathering direct input
- Participation in selected standardisation bodies

UNIPG, CSIC and INAF partners have a great wealth of experience in designing, producing and deploying MPI applications under gLite. These range from relatively simple codes, to large-scale production workflows using multiple externally-provided (and widely used) MPI-enabled libraries. TCD and CSIC will engage with the ARC and UNICORE communities, and will produce high-level documentation for MPI application development and submission under these middleware. In project year one, the first MPI “cookbook” will be produced addressing MPI application development and MPI job submission. This shall be reviewed and updated during project years two and three, as expected middleware changes and new features, such as generic “Parallel” application support, OpenMP support, and GPU application support are introduced. In the past year, we have seen a wide-scale increase in the number of sites supporting GPU clusters. As a result, effort will be required to ensure that the middleware and sites support their integration in a consistent, uniform manner.

A clear weakness in the current MPI software roadmap relates to vendor-supplied distribution of some MPI workbenches versus custom versions that may be provided on a voluntary basis by the EMI project or the grid community. Based on past experiences, many of the vendor-supplied solutions have not been suitable for use by grid sites because they may lack specific features such as Infiniband support or Torque/Maui integration. It is to the project’s benefit to ensure that, where

possible, the vendor solution should be used. This may require greater interaction with the upstream suppliers.

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 first survey will be produced in project month 16.

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

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 MS603 [MS603] for the first 4 months of this Project and the following 12 months will be covered in detail by MS610, due in month 16, together with the next step of this Roadmap, MS612. In the following sections a brief intermediate report and planning is provided.

3.1.1 HammerCloud

HammerCloud [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 – the service is currently being integrated into the daily grid operations for those experiments.

Plans for 2011 include:

- Complete the deployment of HammerCloud for CMS and LHCb
- Develop an extension which detects site failures and automatically excludes sites from the distributed analysis systems
- Develop an alarm system that detects specific site problems by comparing the site metrics with their expected values
- Evaluate new database backends (currently MySQL), including Oracle or a NoSQL solution such as Cassandra

3.1.2 CRAB and CRAB Analysis Server

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

- Develop an asynchronous stage-out tool for user produced data handling. The plan is to couple this tool with CRAB3 server, but keeping it as an independent tool. The new tool will be released by the end of March.
- Develop the web-services based on a REST interface to enable the interactions between the CRAB3 Client and Server. Deliver a prototype by Q2 to be tested by CMS integration. The client prototype will support the 4 basic functionalities: creation, submission, check status, output

retrieval.

- Definition and implementation of a job sandbox distribution strategy including the evaluation and performance study of both HTTPS and GridFTP. The first implementation is due by end of September.
- Consolidation of the client and the RESTful interface, implementing additional functionalities: kill, postMortem. The final client version is planned to be available by the end of September.
- Implement the support for the automatic job resubmission by end 2011. A pluggable agent is needed to offer the possibility of implementing use case specific algorithms for job resubmission

3.1.3 Data Management

3.1.3.1 ATLAS

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 [ACM]. Following this model, ATLAS sites are grouped into ten clouds by geographical and organizational reasons. Each cloud is formed by one Tier1, which must provide a high level of service and is responsible for data storage and reprocessing, and several Tier2s and Tier3s, which are used for analysis and Monte Carlo production and depend directly on the Tier1.

The main directions of work are focused on adapting ATLAS DDM to recent evolutions of the Computing Model:

- ATLAS is driving towards a more dynamic data placement model where the replication of datasets is based on its usage rather than on static pre-placement. For the dynamic data placement model to be optimal, it is convenient to relax the hierarchic boundaries of the grid and allow cross-cloud transfers between Tier 2s. In this respect ATLAS DDM has recently enabled a transfer mechanism between different clouds, which uses channel transfer statistics to optimize these file transfers. The following step is to commission the full mesh of site-to-site links and optimize the source and path selection mechanisms according to the acquired experience.
- Another area of work is the adaptation of DDM Site Services (the set of agents responsible for data placement) to Tier3 analysis sites. While these sites are allowed to have local control of their own resources, it is of general interest to have a unified solution to make the data available to them. As a consequence, DDM Site Services will be extended to respect particular requirements and generate new plug-in libraries that address the data discovery for these particular sites.

In addition to the above, some effort has to be invested in the general consolidation and support of the existing DDM framework, as well as on the improvement of the service level monitoring in order to allow early detection of system failures.

3.1.3.2 LHCb

The DIRAC system [DIRAC] was developed in order to provide a complete solution for using the distributed computing resources of the LHCb experiment at CERN for data production and analysis. One of the most important components is the data management system, whose support in EGI-InSPIRE project has started in October 2010.

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

- Develop and integrate new components to ensure the consistency of LHCb data stored in grid

storage systems and the central file catalogues. This development has been considered necessary on the basis of the experience of managing large volumes of data in the distributed computing environment. Since the storage systems and catalogues are completely decoupled, these checks are non trivial and require the implementation of some dedicated components in order to do the checks in an automated way without requiring any manual operation.

- LHCb Computing Model foresees to keep only the two latest versions of (re)processed real data on disk. However, it is important to keep datasets that were used for a very long time in order to produce published physics results. This requires a tape archiving mechanism, whose implementation is planned for next year.
- Integration of new systems to provide a sustainable usage of grid storage resources, in particular to monitor the usage of datasets over the grid storage systems, provide an efficient policy to remove the least used data, create further replicas of the most used datasets and optimize the use of storage resources.

3.1.4 Persistency and Conditions

The Persistency Framework [PERS-FRAME] is one of the projects set up within the LCG Application Area to provide common software solutions to the LHC experiments at CERN. It consists of three software packages (POOL, CORAL and COOL) that are used by three of the LHC experiments (ATLAS, CMS, LHCb) for storing and accessing several different types of scientific data. CORAL is a generic abstraction layer with an SQL-free API for accessing data stored using relational database technologies. POOL is a generic hybrid technology store for C++ objects and object collections, using a mixture of streaming and relational technologies. COOL provides specific components to handle the time variation and versioning of the conditions data of the LHC experiments.

Plans for 2011 include:

- User support, service operation and software maintenance. These tasks require a large support and development effort as they include, for instance, the preparation of regular software releases (one per month on average, following explicit requests of the experiments for urgent bug fixes, functionality enhancements or upgrades in external software dependencies such as Boost or ROOT) and the follow-up of service operation incidents (which normally result in bug fixes in the Persistency Framework code, but often end up in the need for a more global troubleshooting, involving for instance the Oracle server software or the grid middleware).
- Development of new functionalities such as the capability of the CORAL software to restore the database connection and session after a network glitch.
- R&D to evaluate new technologies relevant to data access optimization.

3.2 Services for the Life Science Virtual Research Community

To coordinate their efforts and sustain their activity, members from the Life Science community self-organized into the project-independent “Life Sciences Grid Community” (LSGC) over the first period of the EGI-InsPIRE project. The LSGC is currently representing 3 Virtual Organizations (biomed, lsgrid, and vlemed). It receives support from 6 NGIs (Dutch, French, German, Italian, Spanish and Swiss NGIs) and the HealthGrid association[1].

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

As scheduled in MS602 [MS602], in 2010 the LSGC described and agreed on internal governance rules[2]. Developed internal communication channels (monthly phone conference, mailing list, wiki[3]). Setted up a technical team monitoring the infrastructure through a VO-specific Nagios server[4] and tracking the issues addressed to the community through the GGUS ticketing system.

The LSGC is also currently designing a user management database, which will facilitate liaising with hundreds of users registered in the affiliated virtual organizations. The development roadmap of this user database was introduced in MS602 [MS602]. The following steps have been completed in 2010:

- Design of the LS VRC user database
- Analysis of the "VO admin tool" capability
- Implementation of the file access control to the VO LFC (notion of user home directory)

The development of the user database, initially scheduled from October to December 2010, was delayed as the LSGC effort was focused on the set up of the community and there appeared to be development dependencies with external tools (VOMS, application database) where the available interface still needs to be specified. The development tasks are therefore rescheduled as follows:

- January 2011 (M9): interaction with VOMS and application database developers to specify available interfaces
- March 2010 (M11): implementation of the LS VRC user database and its web front-end; link with the applications database.
- May 2010 (M13): first release of the tools supporting VO membership management; automatic updates of the VO mailing lists from the DB and tools.
- September 2011 (M17): second release of the tools including the implementation of the group (projects, national…) management procedures.
- December 2011 (M20): third release of the tools including updates based on usage and evolutions
- June 2012 (M26): fourth and final release of the tools.

In addition, the user database management tools will be interfaced to the LSGC Dashboard.

3.3 Services for the A&A Virtual Research Community

The main areas of interest for the A&A community are related to Visualisation Interface to the Virtual Observatory (VisIVO) [VISIVO] and other visualization tools, grid and supercomputing with a special interest for the use of MPI in grid DCIs and the databases. Each of these areas is covered by a specific sub-task of TSA3.5.

3.3.1 VisIVO & Visualization Tools

3.3.1.1 Current Status

VisIVO is a suite of software tools aimed at creating 3D customized views of many type of wide used datasets. Its peculiar characteristic is that there is no limit for what concerns the size of input tables

[2] LGC purpose and rules document, https://dav.healthgrid.org/lsvrc/LSVRC_proposition_09-08-2010-final.pdf

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

[4] LSGC nagios monitoring server, https://grid04.lal.in2p3.fr/nagios/

containing data to be processed, thus they are able to support very large scale datasets (tens of Terabytes and more). The linear scaling observed in our tests suggests that handling large datasets is effectively restricted only by the underlying file system limitations.

In order to foster a wide adoption of the tool it has been designed and developed to be compliant with different kind of hardware/software platforms.

The current release of the VisIVO Server can run on many hardware architectures with different Operating Systems: different flavours of Linux, Mac and Windows; no specific additional resources are required to the Operating System.

The gridified version of VisIVO server is being developed so that it will not depend on a particular implementation of the Grid middleware.

The VisIVO Server consists of three core components: VisIVO Importer, VisIVO Filter and VisIVO Viewer respectively. To create customized views of 3D renderings from astrophysical data tables, a two-stage process is needed. First, VisIVO Importer is used to convert user datasets into VisIVO Binary Tables (VBTs). Then, VisIVO Viewer is invoked to display customized views of 3D renderings. VisIVO Filters are collections of data processing modules able to explore datasets enhancing and highlighting their hidden properties.

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

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

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

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

3.3.1.2 Plans for the next six months

Plans for VisIVO have been revised and modified with respect to the previous issue of this document (MS602 [MS602]). During 2011 the following steps will be accomplished:

a) the VisIVO Libraries API (Application Programming Interface) will be completed; some use cases will be prepared and tested on the gLite infrastructure
b) prototypes for importers and filters will be implemented and tested
c) a MPI-compliant version of VisIVO will be deployed.

VisIVO Libraries represent a very important step in order to export the usage of the 3D visualization rendering and movie creation tools to other user communities. When this phase will successfully

end, the MPI-compliant version of VisIVO will be tested on the infrastructure, thereafter it will be deployed.

Concerning the next six months, from Month 11 (March 2011) to Month 16 (August 2011) the following milestones have been scheduled:

- Month 12 (April 2011): VisIVO API deployment. The grid-enabled version of VisIVO API with some use-cases will be deployed. The results of some basic tests will be published.
- Month 14 (June 2011): Prototypes for Importers and Filters implemented. The first tests will be done.

Milestone	M	Title	Description
MSA3.5.2.1	12	VisIVO API prototype	Grid enabled version of VisIVO API with some use-cases deployed
MSA3.5.2.2	14	Importers and Filters prototypes	Prototypes for Importers and Filters implemented. First tests done
MSA3.5.2.3	18	VisIVO-MPI version	MPI-compliant version of VisIVO deployed
MSA3.5.2.4	22	Prototype of VisIVO for CUDA	First prototype of VisIVO for CUDA generated through GPUs on the worker nodes
MSA3.5.2.5	25	Final version of VisIVO for CUDA.	Final version of VisIVO for CUDA generated through GPUs on the worker nodes
MSA3.5.2.6	31	I/O in Grid catalogue	Final version of VisIVO allowing I/O operation directly on the Grid Catalogue
MSA3.5.2.7	36	Portal for VisIVO	Final version of a portal that allows the usage of VisIVO on the grid.

3.3.2 Grid and Supercomputing

3.3.2.1 Current Status

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

Because cosmological simulations represent one of the most important classes of A&A applications requiring HPC resources, we identified the following applications: FLY (INAF-OACT Cosmological code) and Gadget and Flash, the most common cosmological codes in Astrophysics.

We are now in the process of collecting requirements from these applications following this schema:

a) Preparation of the initial dataset; its size is of several hundreds of Gigabytes.
b) Data production phase, generally performed through parallel code whose execution involves hundreds of CPU/cores. We are now starting the design of some preliminary tests to run in gLite.

Another activity that is worth to mention relates the MPI-WG; members of the A&A community are directly involved in this working group.

3.3.2.2 Plans for the next six months

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

Concerning the next six months, from Month 11 (March 2011) to Month 16 (August 2011) the following milestones have been scheduled:

- Month 12 - April 2011 - Design of a number of Use-Cases and complex workflows encompassing the complete set of Use-Cases - As stated above, the collection of a number of use-cases and complex workflows to demonstrate the importance of combining HPC and grid resources for A&A research is a continuous process started in the early months of EGI-InSPIRE. Most of this material has been selected in the last months; it will be now gathered in a unique document to be delivered at MSA3.5.3.1.
- Month 15 - July 2011 - Definition of those indicators suitable to gauge the usage of HPC resources in grid - After the issue of the document scheduled at MSA3.5.3.1, it is necessary to prepare a validation plan against which the implemented solutions will be tested. In practice, the Validation Plan will define a number of indicators whose values will allow the goodness of the proposed solutions to be gauged. The validation plan will be detailed in the document issued at MSA3.5.3.2.

Milestone	M	Title	Description
MSA3.5.3.1	12	Requirements, Use-Cases and Workflows defined	Design of a number of Use-Cases and a complex workflow encompassing the complete set of use-cases
MSA3.5.3.2	15	Validation Plan in place	Definition of the indicators allowing to gauge the usage of HPC resources in Grid
MSA3.5.3.3	36	Global Test and results analysis completed	Complete test involving all workflows: medium and large size

3.3.3 Databases

3.3.3.1 Current Status

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

During this project period the following milestones have been scheduled:

- Month 10 - February 2011 - Evaluation of tools and services - A document will be issued at MSA3.5.4.1 where all evaluated tools and services will be illustrated. Those of them selected to be used for A&A applications will be listed. The document will also expose the selection criteria and the motivations for the adopted choices.

Milestone	M	Title	Description
MSA3.5.4.1	10	Tools and services evaluated	Tools and services evaluated. Specific activity planned

3.3.3.2 Plans for the next six months

Integration between grid and databases remains one of the hot topics for A&A given that astronomical databases are one of the key resources used within the A&A HUC, especially in the context of the Virtual Observatory. Without a good support for databases the grid does not have chances to penetrate the A&A community. The evaluation of tools and services, like AMGA, GRelC and possibly others, will continue.

New activities will be jointly carried out with EGI.eu and with NGIs if some of the evaluated tools/services need to be modified to fully meet the requirements of the A&A applications.

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

The plan for the forthcoming months includes a complete characterization of the test for FLY, Gadget and Flash software codes; the test will be run in gLite. We will also identify a typical workflow, to be executed on the grid, encompassing the following steps:

a) Preparation of the initial dataset
b) Production of data
c) Post-analysis phase
d) Persistency of produced data (the integration between grid and database and any other related aspect will be handled).

Concerning the next six months, from Month 11 (March 2011) to Month 16 (August 2011) the schedule (as shown in table below) does not foresee any specific milestone for this subtask.

Milestone	M	Title	Description
MSA3.5.4.2	22	Mid-term release	Mid-term release of new/customized tools and services
MSA3.5.4.3	36	Final release	Final release of new/customized tools and services

3.4 Services for the Earth Science Virtual Research Community

The ES HUC is concentrated in this work package on the access to ES Data infrastructures (1-3). But the ES HUC of this WP does also unfunded operating support for ES and raises new topics of the community that are of increasing interest for ES EGI users (see 4-5).

1. From the beginning of the project the activity was concentrated on interfacing to the GENESIS-DEC project (mainly ESA satellite data). In the last period new features that are under development in the on-going GENESI-DEC project were discussed and the the existing solution was maintained for interested users. This will be continued in the next period and new features will be implemented as soon as possible which are required by ES HUC users.
2. A closer collaboration with the ES Community in the French NGI is established. Together with the Institute Pierre Simon Laplace (IPSL) in Paris the interfacing of the EGI infrastructure with the Earth System Grid (ESG) will be started. It is planned to access a CMIP5 Data Node from EGI. The underlying technology here is based on OPenNDAP (see MS602 3.4). This project will be done in collaboration with other climate projects in France and in Europe. The work will be partly supported by unfunded manpower.
3. In the next period the transformation of the GRelC testbed, originally introduced for climate data, to a production environment for other ES applications will take place. The service itself and the maintenance is part of 2.3.2.
4. The operation and maintenance of a “catch all ES VO” will be continued unfunded.
5. The ES HUC has started to establish a network inside the HUC to exchange the knowledge about porting of ES simulation codes on new Hardware Architectures on Many- (GPU) and Multicore. There is an increasing interest in the ES community. ES disciplines like commercial seismology (e.g. CGG-Veritas, former ES HUC partner) introduced GPUs very early. The activity is not part of the work package directly and remains unfunded.

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 are 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. 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 D6.1 “Capabilities Offered by the HUCs to Other Communities” [D6.1], as well as MS601 “HUC Contact Points and Support Model” [MS601]. The achievements of the workpackage 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 workpackage, 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 (to appear)
MS608	MS608 – Integration of the VisIVO server with the production infrastructure: https://documents.egi.eu/document/328 (to appear)
D6.1	D6.1 – Capabilities Offered by the Heavy User Communities: https://documents.egi.eu/public/ShowDocument?docid=154
D6.2	D6.2 – Sustainability plans for the activities of the Heavy User Communities: https://documents.egi.eu/public/ShowDocument?docid=309
D6.3	D6.3 – Annual Report on the Tools and Services of the Heavy User Communities: https://documents.egi.eu/document/312
MONARC	MONARC Technical Notes: http://www.cern.ch/MsONARC/docs/monarc_docs.html
DASHBOARD	Experiment Dashboard for Monitoring of the LHC Distributed Computing Systems; J. Andreeva et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
GANGA	Ganga: a tool for computational-task management and easy access to Grid resources; *Computer Physics Communications, Volume 180, Issue 11, (2009)*

GRELC	The GRelC Project from 2001 to 2011, ten years working on Grid-DBMSs; S. Fiore et al, to appear in book of "Grid and Cloud Database Management", Springer, 2011
KEPLER	The Kepler Workflow Tool, available online: http://kepler-project.org
GRIDWAY	The GridWay Framework for Adaptive Scheduling and Execution on Grids Scalable Computing; E. Huedo, R.S. Montero and I.M. Llorente, Practice and Experience 6, 1–8 (2005)
SOMA2	SOMA2 – Open Source Framework for Molecular Modelling Workflows; Tapani Kinnunen, TH Nyrönen and P Lehtovuori, Chemistry Central Journal, Volume 2, Supplement 1, P4, DOI: 10.1186/1752-153X-2-S1-P4 (2008)
HAMMERCLOUD	HammerCloud: A Stress Testing System for Distributed Analysis; Daniel C. van der Ster et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
WIKI	https://wiki.egi.eu/wiki/SA3:_Services_for_the_Heavy_User_Community
CRAB	CRAB: A CMS Application for Distributed Analysis Nuclear Science, IEEE Transactions on Volume 56, Issue, Part 2, Oct. 2009 pp:2850 - 2858
CMSDA	Distributed Analysis in CMS; Journal of Grid Computing, Vol.8, Number 2, 159-179, (2010) doi: 10.1007/s10723-010-9152-1
CTDR	CMS Computing: Technical Design Report; M. Della Negra et al, CERN-LHCC-2005-023
ACM	The ATLAS Computing Model; R. Jones and D. Barberis, CERN-LHCC-2004-037-G-085
DIRAC	DIRAC – The Distributed MC Production and Analysis for LHCb; A. Tsaregorodstev et al, Computing in High Energy and Nuclear Physics (CHEP'04), Interlaken, Switzerland (2004)
PERS-FRAME	LCG Persistency Framework (POOL. CORAL, COOL): status and outlook; A. Valassi et al, Computing in High Energy and Nuclear Physics (CHEP'10), Taipei, Taiwan (2010)
VISIVO	VisIVO: Data Exploration of Complex Data; G. Caniglia et al, Mem. S.A.It. Vol. 80,441 (2009)