### 1.5.5 Work Package SA3: Services for Heavy User Communities

#### 1.5.5.1 Summary

| | | | | | | | |
|---|---|---|---|---|---|---|---|
| **Work package number** | SA3 | **Start date or starting event:** | | 1/5/2010 | | | |
| **Work package title** | Services for Heavy User Communities | | | | | | |
| **Activity type** | SVC | | | | | | |
| **Participant number** | 10 | 12 | 13 | 14 | 19 | 21 | 28 |
| **Participant short name** | KIT | CSIC | CSC | CNRS | TCD | INFN | CYFRO NET |
| **Person-months per participant** | 27 | 45 | 18 | 83 | 18 | 125 | 6 |
| | | | | | | | |
| **Participant number** | 31 | 32 | 35 | 37 | | | |
| **Participant short name** | ARNES | IISAS | CERN | EMBL | TOTAL | | |
| **Person-months per participant** | 3 | 18 | 341 | 45 | 727 | | |

**Objectives**

This activity provides continued support for activities currently supported by EGEE while they transition to a sustainable support model within their own community or within the production infrastructure by:

- Supporting the tools, services and capabilities required by different heavy user communities (HUCs)
- Identifying the tools, services and capabilities currently used by the HUCs that can benefit all user communities and to promote their adoption
- Migrating the tools, services and capabilities that could benefit all user communities into a sustainable support model as part of the core EGI infrastructure
- Establishing e a sustainable support model for the tools, services and capabilities that will remain relevant to single HUCs

**Description of work**

**TSA3.1: Activity Management**

The SA3 management is the responsibility of the Activity Manager, who reports to the UCO for the running of this activity and coordinates the work across the different tasks within the activity to meet the stated objectives. This activity will be represented in the USAG by the Activity Manager and the task leaders for the sub tasks in TSA3.2 to ensure that the services being offered to the general EGI community are meeting their needs. The domain specific activities will have their community specific mechanisms for supervision. The Activity Manager will pay particular attention to the provision of the shared services and tools task (TSA3.2) to ensure that all requests from all communities are correctly evaluated and prioritised into the work plan even if the work is being undertaken by a single community. All tasks will establish outreach and sustainability plans, and mechanisms for monitoring and gathering feedback on the delivery of their services.

The work in SA3 has many benefits to other HUCs (e.g. EIRO and ESFRI) and more broadly the general DCI community. Some of the potential results from this activity include: dashboards customised to specific VOs, workflows and schedulers bridging different DCIs (e.g. DEISA, PRACE, EGI), support of MPI, frameworks for managing collections of jobs on DCIs, services for accessing relational data resources, secure data storage, visualisations tools, etc. Technical dissemination material will be generated by SA3 to inform all communities of this work.
Requested Effort: 18PMs CERN

**TSA3.2: Shared services and tools**

The services and tools within this task are already being used by multiple communities, and have been identified as having benefit to other specific communities or the general user community. It is expected that by the end of the project many of these services and tools will have been migrated into UMD (i.e. be supported independently of the EGI project by an external software provider) and deployed within EGI as part of the core, supported infrastructure. This migration will be justified by the continued adoption and use of these tools and services by HUCs and the broader EGI user community.

**Subtasks**:

- **TSA3.2.1 Dashboards**: Dashboards provide a generic framework to monitor sites and their services within a VO using tests specific to that community. Dashboards have emerged from within the HEP community, and are now being adopted by the LS community, to monitor their resources. Requested Effort: 60PMs CERN, 12PMs CNRS
- **TSA3.2.2 Applications**: The GANGA and DIANE tools are both part of the EGEE RESPECT programme which recognises software that builds on top of the gLite platform. Although initially developed for the HEP community these tools have now gained traction in other communities, as they provide simple environments to manage large collections of tasks. The requirements of these tools will be integrated into the workplan. Requested Effort: 60PMs CERN, 18PMs UI SAV, 3PMs ARNES
- **TSA3.2.3 Services:** HYDRA and GReIC are services that have emerged from a single community that show potential for adoption in others. HYRDA allows an encryption key to be securely stored on distributed servers in order that storage elements can be used to store confidential data which is critical for the medical community securely. The GReIC service provides uniform relational and non-relational access to heterogeneous data sources and is currently being used to support bioinformatics and Earth Observation Systems. Requested Effort: 18PMs CNRS, 27PMs INFN
- **TSA3.2.4 Workflow and Schedulers**: These tools are critical in integrating complex processes, generally involving multiple data sources and different computational resources, as needed within many disciplines. SOMA2 is a web-based workflow tool used for computational drug design and general molecular modelling. TAVERNA is used extensively by the bioinformatics community. The combination of the Kepler workflow engine and the Migrating Desktop platform are used by the Fusion community to run workflows requiring visualisation and interactive access on gLite and UNICORE-enabled resources. For simpler workflows and metascheduling scenarios the GridWay system is used by the Fusion community. Effort is provided to maintain the integration of these tools with the different systems. Requested Effort: 27PMs CSIC, 18PMs CSC, 6PMs CYFRONET, 18PMs EMBL
- **TSA3.2.5 MPI:** Support for parallel computing (MPI) applications are critical for many user communities but the integration of this capability into the general infrastructure has been difficult. This task will focus on the improvement of the core services and software needed to support MPI, while engaging with two representative user communities (CCMST & Fusion) to ensure that the offered support meets their requirements. Requested Effort: 27PMs CSIC, 18PMs TCD, 9PMs INFN

**TSA3.3: Services for High Energy Physics (HEP)**

The HEP VO specific services are devoted to the support of the Grid interfaces of the 4 LHC experiments and are of particular importance now as we enter the exploitation phase of the world's largest scientific machine – the Large Hadron Collider at CERN. They will ensure the LHC experiments can rely on the Grid for their data handling, as planned in the LCG Technical Design Report[11].

All LHC VOs require optimizations and improvements in a number of common areas: workload management, data management, monitoring, service deployment and operation. These optimisations – of which specific details are provided in the task description below – are an essential part of the move to sustainable operations, as well as handling additional load and complexity expecting from LHC data acquisition, (re-)processing and analysis. An agile structure will be used to organise this work allowing the team to react quickly to high-priority needs of the LHC VOs as experience of the accelerator grows. Requested Effort: 204PMs CERN, 60PMs INFN

**TSA3.4: Services for Life Sciences (LS)**

This task will provide services and service deployment for the Life Sciences community. The Medical Data Manager (MDM) will be used to store image data on services attached to the Grid. The resources available to the biomolecular community at the EBI will be integrated into EGI. The WISDOM Production Environment (WPE) will be evolved to integrate new middlewares (ARC, UNICORE) at job and data management levels

[11] http://lcg.web.cern.ch/LCG/TDR/LCG_TDR_v1_04.pdf

Requested Effort: 54PMs CNRS, 27PMs EMBL

**TSA3.5: Services for Astronomy and Astrophysics (A&A)**
The A&A community will Grid enable visualisation and associated data interpretation tools (e.g. VisIVO) to show how data, that will in future be collected from projects such as the Sloan Disgital Sky Survey, the Low Frequency Array (LOFAR), the Square Kilometre Array (SKA) and others, will be analysed using the production infrastructure to rapidly provide visual representations of this data. HPC (e.g. DEISA/PRACE) resources and HTC resources will be integrated in this task as part of this proof of concept.
Requested Effort: 30PMs INFN

**TSA3.6: Services for Earth Sciences (ES)**
Implement, deploy and maintain the EGDR service to provide access from the grid to resources within GENESI-DR.
Requested Effort: 27PMs KIT

**Deliverables** (brief description) and month of delivery

**Month 3, 15 & 27: DSA3.1.1-3 – Capabilities offered by the HUCs to other communities**
This public report will illustrate how the functional capabilities being supported with this activity could be reused by other communities using European DCIs. Sufficient technical depth should be provided for potential adopters of DCI platforms to make an initial assessment of how they could work with the offered technologies.

**Month 9, 21 & 33: DSA3.2.1-3 – Sustainability plans for the HUC activities**
The activities taking place within SA3 must plan for a sustainable future – either within their own community, as part of the generic production infrastructure, or through some other mechanism. The sustainability plan for each activity within SA3 will be defined in this public report by assessing the progress made to date and providing plans for the next year.

**Month 11, 23 & 35: DSA3.3.1-3 – Annual report on the HUC Tools and Services**
A public report describing the current status of the tools and services provided within SA3, their activity over the preceding year, and the future activities including measures to assure sustainability of the current work.

#### 1.5.5.2 Overview

This activity responds to the 1.2.1.2 sub-call, whose objective is *"to deploy services for user communities that are heavy users of DCIs and have a multi-national dimension. Software components should be integrated in platforms as needed for service provision. Where appropriate, new service provision models should be explored and harmonised interfaces to DCI resources should be ensured."*

This activity has three main goals:

- 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 through 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 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. This is especially critical for communities such as those that are already actively exploiting the infrastructure.

These communities have been identified through their current usage of the EGEE and related infrastructures (e.g. EUFORIA, EDGES, national infrastructures) by collecting the usage statistics through the relevant accounting mechanisms, and are shown in Figure 13 below (1 CPU approximated to be 1000 SpecInt2000).

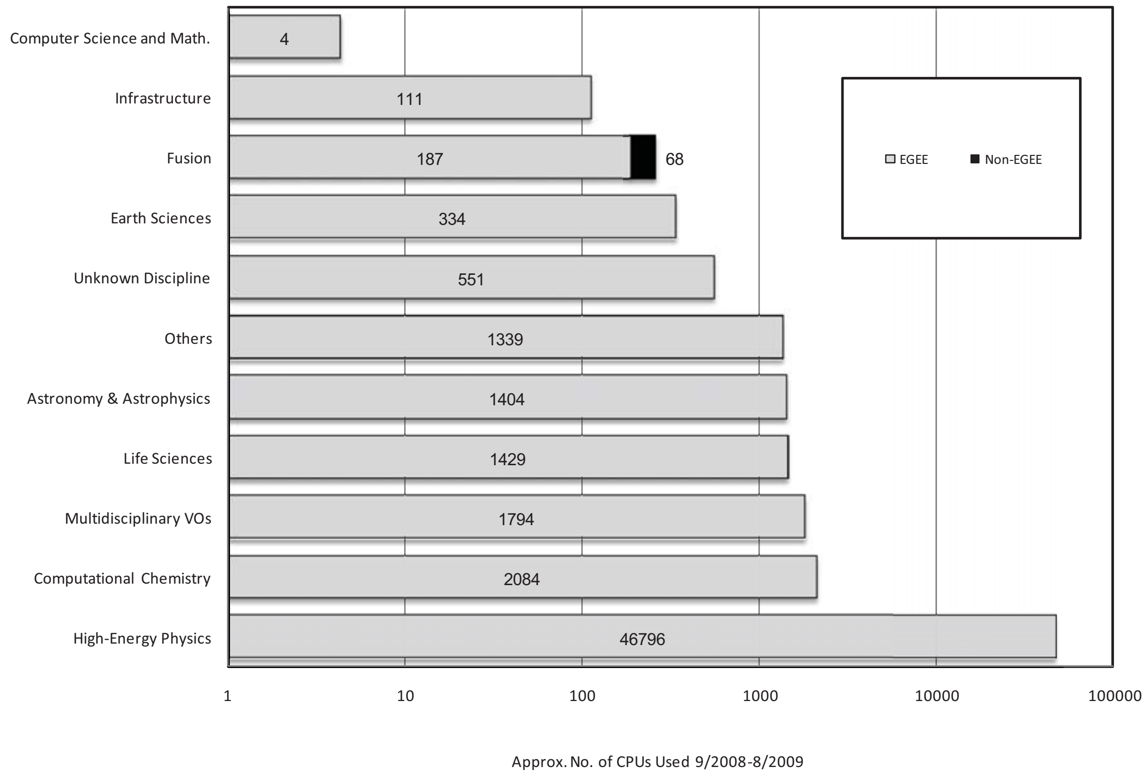


**Figure 13: User Communities resource usage**

The communities identified as Heavy Users Communities (HUCs) within this proposal are:

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

Besides their extensive usage of grid infrastructure, these communities, and especially HEP and LS, have played a central role in EGEE, actively working to bring the grid to production quality. Their continual feedback on the deployed infrastructure (the effectiveness of the services, and their functionality, stress tests, operational procedures, etc.) needs to be captured within EGI, and from all the HUCs in order to develop EGI's service offering for all user communities.

The tools, services and procedures selected for support are clearly essential to the communities that are already using them, but the supported activities also show significant potential for uptake by other communities as they can all accelerate a community's adoption and effectiveness when using DCIs. The supported activities include:

- Task Management: For many communities, a job can be broken down into a large number of independent tasks (e.g. parameter studies). Managing these collections of tasks, or indeed several collections simultaneously through manual means is very complex and prone to errors and omissions if any of the jobs fail. The GANGA and DIANE packages address such scenarios and can benefit any discipline that has similar analysis patterns.
- Workflow: Either within a resource, or between several resources, many jobs break down into sets of inter-dependent tasks. Workflow engines have become a common tool in allowing an end-user to manage their complex use of resources. The TAVERNA and GridWay environments include powerful workflow and resource selection engines that simplify coordinated access to distributed resources, while SOMA2 provides workflows through a web browser.
- MPI: The Message Passing Interface is a common parallel programming model that is essential for large-scale scientific modelling on clustered resources. Its availability on the production infrastructure is

becoming critical for some user communities, and its routine and reliable availability would make use of these resources attractive for many others.

- Dashboards: Monitoring of the core infrastructure is assured through the operations activity (SA1). However, many communities need to be assured that specific functionality and services unique to their community are operating effectively. The Dashboards task provides a generic framework that could be customised by any community. The effectiveness of the developed interfaces to customisation will be demonstrated through the HEP and t LS communities.
- Data Access: High throughput data analysis is one of the key use cases for the production infrastructure. Increasingly, this requirement is moving beyond file-based data to relational, structured or image-based data. Several tasks in this activity work to improve data access by providing relational access to legacy databases (GReIC), making new data sets available from the production infrastructure (Biomolecular data sets), providing access to medical image data (Medical Data Manager), and ensuring that data can be securely encrypted for confidential data storage (Hydra).
- Visualisation: For large scientific data sets the only viable means to comprehend the collected data is through visualisation. However, for large data sets the visualisation capability also needs to be grid enabled. The integration of VisIVO with the production resources (TSA3.6) will benefit the any scientific communities that need to process and visualise large datasets

Together, these capabilities provide an attractive offering for many new communities. Their development and support by these HUCs will be supported by the

- Maintenance, operation and development of specific software services, tools and applications.
- Deployment and integration of domain-specific applications and tools with the infrastructure
- Integration of data resources into the e-infrastructure
- Expansion of the user base beyond the initial user community

Some of these tasks are targeted at specific user communities, while others are driven by a subset of the HUCs for the benefit of all current users and by the integration of this work into the core infrastructure provided for all users.

#### 1.5.5.3 TSA3.1 Activity Management

The SA3 activity will be managed by CERN. The SA3 Activity Manager will report to the UCO and be a member of the OMB to ensure that these services are integrated into EGI's generic service offering within SA1 and the USAG. The USAG will be one source of feedback outside of the specific user communities covered by this activity and will provide opportunities for this work to be integration with the EGI support process. The Activity Manager will also drive within each task the production and maintenance of technical dissemination material. This work will be used as the basis for material illustrating how the adoption of DCI technologies could evolve the work of applied researchers.

#### 1.5.5.4 TSA3.2 Shared Services & Tools

The work around each tool or service is led by one or more HUCs funded within this activity, while the development of community specific plug-ins and any associated prototyping are outside scope of this activity and will be carried out where needed within the VRCs or the broader user community. Support processes will be defined for each activity to ensure that requirements from all communities (including the funded communities) are treated equally. The Shared Services & Tools supported in this task are detailed in Table 13 below – along with the funded, current and future HUC user community.

**Table 13: Supported shared tools and services**

| Item | Functionality | Funded HUC | Current HUC Adopters | Future HUC Adopters |
|---|---|---|---|---|
| Dashboard | Monitoring | HEP | HEP, LS, CCMST, ES | A&A |
| GANGA | Task Management | HEP | HEP, CCMST, A&A, ES | F, LS |
| DIANE | Task Management | A&A | A&A, ES, LS | |
| GRelC | Database Access Service | ES | ES, A&A, CCMST | LS |
| Hydra | Encrypted Data | LS | LS | |
| SOMA2 | Web based workflow | CCMST | LS, CCMST | |
| TAVERNA | Workflow Manager | LS | LS, A&A, CCMST | ES |

| | | | | |
|---|---|---|---|---|
| Roaming Access Server (RAS) & Migrating Desktop (MD) | Running Interactive Applications | F | F, ES | A&A, LS |
| Gridway | Resource Broker | F | F, ES | LS |
| MPI | Programming Model | CCMST & F | CCMST, F, ES | A&A |

A brief description of each sub-task follows.

#### 1.5.5.4.1 TSA3.2.1 Dashboards

In order to perform production and analysis tasks across a highly distributed system crossing multiple management domains, powerful and flexible monitoring systems are clearly needed. The experimental Dashboard monitoring system was originally developed to support the four main LHC experiments (LHCb, CMS, ATLAS, ALICE). This framework not only supports multiple grids / middleware stacks, including gLite (EGEE), VDT (OSG) and ARC (NDGF), but is also sufficiently generic to address the needs of multiple other communities including, but not limited to HUCs. Furthermore, this framework covers the full range of the experiments' computing activities: job monitoring, data transfer and site commissioning. It also addresses the needs of different categories of users, including:

- Computing teams of the LHC VOs;
- VO and WLCG management;
- Site administrators and VO support at the sites;
- Users running their computational tasks on the grid infrastructure.

Future work will concentrate effort on common applications which are shared by multiple LHC VOs, but which could also be used outside the scope of LHC and HEP. Examples of such applications are: generic job monitoring and user task monitoring, FTS monitoring, site status board, and VO-specific site availability based on the results of tests submitted via the Site Availability Monitor (SAM).

Reliable monitoring is a necessary condition for establishing and maintaining production quality of the distributed infrastructure. Additionally, monitoring the computing activities of the main communities using this infrastructure provides the best estimation of its reliability and performance. The importance of flexible monitoring tools focusing on applications has been demonstrated to be essential not only for "power-users", but also for single users. For the power users (such as managers of key activities such as large simulation campaigns in HEP or drug searches in LS communities) a very important feature is to be able to monitor the resource behaviour to detect the origin of failures and optimise their system. They also benefit from the possibility to "measure" efficiency and evaluate the quality of the service provided by the infrastructure. Single users are typically scientists using the Grid for data analysis, and verifying hypotheses on data sets that could not be achieved on other computing platforms. In this case the monitoring/dashboard is a guide to understanding the progress of their activity, and identifying and solving problems connected to the application.

It is essential to allow efficient user support by "empowering the users" in such a way that only non-trivial issues are escalated to support teams (for example, jobs on hold due to scheduled site maintenance can be identified as such and the user can decide to wait or to resubmit).

#### 1.5.5.4.2 TSA3.2.2 Applications

**GANGA** is an easy-to-use frontend for job definition and management, implemented in Python. It has been developed to meet the needs of ATLAS and LHCb for a Grid user interface, and includes built-in support for configuring and running applications based on the Gaudi / Athena framework common to the two experiments. GANGA allows trivial switching between testing on a local batch system and large-scale processing on Grid resources.

A job in GANGA is constructed from a set of building blocks. All jobs must specify the software to be run (application) and the processing system (backend) to be used. Many jobs will specify an input dataset to be read and an output dataset to be produced. Optionally, a job may also define functions (splitters and mergers) for dividing a job into sub-jobs that can be processed in parallel, and for combining the resultant outputs. GANGA provides a framework for handling different types of application, backend, dataset, splitter and merger, implemented as plugin classes. Each of these has its own schema, which displays the configurable properties.

As it is based on a plugin system, GANGA is readily extended and customised to meet the needs of different user communities. Activities outside of ATLAS and LHCb where GANGA is successfully used include GÉANT4 regression tests and image classification for web-based searches. GANGA is included in the EGEE RESPECT Programme (Recommended External Software Packages for EGEE Communities).

The number of GANGA users has steadily increased and today there are several hundred grid users using the tool in their daily work, some 25% of whom are not from HEP VOs. These VOs have successfully "gridified" applications in a wide range of fields including Fusion, Material Sciences, Accelerator Studies and Biomedical applications. The effort requested here will maintain a production service to the HUCs in the critical years following EGEE, as support transitions to a sustainable model supported by the community.

**DIANE** is a lightweight distributed framework for parallel scientific applications in master-worker model. It assumes that a job will be split into a number of independent tasks which is a typical case in many scientific applications. The DIANE framework takes care of all synchronization, communication and workflow management details on behalf of the application. DIANE is included in the EGEE RESPECT Programme.

#### 1.5.5.4.3 TSA3.2.3 Services

**GRelC (**Grid Relational Catalogue) is a Grid database access and integration service. The GRelC service allows users to interact with different Database management systems, both relational (PostgreSQL, MySQL, Oracle, DB2, SQLite, etc) and non-relational (eXist, XIndice, XML flat files). It provides a uniform access interface to heterogeneous data sources in a grid environment. The GRelC middleware has been included in the EGEE RESPECT Programme since it works well with the EGEE software by expanding the functionality of the grid infrastructure (with regards to database management in the grid). The GRelC middleware is currently used within several grid research projects to support bioinformatics experiments on distributed and huge databases as well as the metadata management related to Earth Observation System applications (i.e. Climate-G).

During the project, the GRelC system (the P2P network of GRelC services deployed within EGI) will be enhanced to support the EGI communities with a new set of functionalities. These will be accessed by end-users through the GRelC Portal, a seamless, ubiquitous and web-based environment for the management of geographically spread and heterogeneous grid data sources.

An important task will be related to the monitoring and control functionalities connected with the underlying P2P infrastructure of the GRelC system. Such a management framework will be managed through the GRelC Portal by means of a new set of web pages exploiting the dashboard approach (charts, reports, table, diagrams able to provide global and local views about the status of the system). Users will be able to configure, manage and query their own GRelC services, exploiting a wide set of management functionalities embedded into the GRelC Portal. A key point will be to make easier (few steps in a web-based wizard) the *gridification process* (bringing into the grid) of a database resource.

A key task will be the creation of the EGI Database of Databases, a registry service accessible through a specific GRelC Portal web page that will contain all of the information about the grid-databases available in the GRelC System. Users will be able to:

- *query the registry* (exploiting a keyword-based approach) asking for specific databases, filtering by VO, keywords, domain, etc. This will help people working in a specific domain to quickly identify available and related resources, identify key people working on specific subjects, easily contact them to establish collaborations, etc.
- *join a specific grid-database*, submitting via the web a request to the grid-database administrator to know more about the supported VOs, etc.;
- *add comments* on the available data and the related data sources, being part of a community exploiting a collaborative and Web2.0 oriented approach. All of this data will be available for future users, creating a knowledge base centred around community-oriented topics.

The EGI Database of Databases will complement the functionalities provided by the EGI Application Database (TNA3.4) and will represent a distributed and multi-VO system.

During the project a set of specific use cases, oriented to the EGI VOs working with the GRelC service, will be defined starting from user needs and requirements. Support will be provided to these VOs to 'gridify' their data sources and to use these experiences to drive the design and implementation of new functionalities provided through the GRelC Portal. Success stories relating to the GRelC Portal will be disseminated through NA2 to the entire community.

**HYDRA** will be maintained, tested and deployed by NGI France (CNRS partner). The Hydra service is a critical component for the MDM service in particular and medical image manipulation on the grid in general. Preserving compatibility of the MDM with the evolution of Hydra and guaranteeing proper operation of the service is critical for the medical imaging user community. This work will cover (1) functionality testing and stress testing of the Hydra service; (2) transparent interface and update of the Hydra encryption / decryption functionality inside the MDM client; (3) deployment and maintenance of a multi-servers Hydra services to serve the community.

#### 1.5.5.4.4 TSA3.2.4 Workflows and Schedulers

**SOMA2** is a web browser-based workflow environment for computational drug design and general molecular modelling (http://www.csc.fi/soma) in the LS community. The purpose of the SOMA2 environment is to provide users with easy access to computational tools. SOMA2 hides all technicalities related to execution of scientific applications in complex computing facilities allowing users to focus on their actual scientific tasks. The SOMA2 server platform will be set up and configured with suitable authentication and user management systems needed for integrating SOMA2 with the EGI infrastructure. Grid job submission features will be integrated and deployed into SOMA2 including procedures for handling personal grid certificates for authentication. These features will be tested and documented to ensure that the features work as advertised. The SOMA2 scientific gateway will operate during the project including the necessary user management, system administration, and maintenance of the existing scientific applications integrated into the gateway and of the new tools and applications requested by the use community.

**TAVERNA** is a workflow engine initially developed within the MyGrid project by University of Manchester for the LS community. Because of its wide adoption by the bioinformatics community, TAVERNA has been ported to EGEE by EGEE-III Life Sciences cluster. TAVERNA is maintained and under further development by a consortium involving most notably the University of Manchester and the EBI. This work will ensure its continual interoperability with the security models used by the production infrastructure and the middleware deployed within EGI (e.g. gLite, ARC, UNICORE).

**MD / RAS / Kepler** platform consists of two major components: server (Roaming Access Server - RAS) and client (Migrating Desktop - MD) that is used by the Fusion and Earth Sciences communities. Users authenticate and login in the Java based graphical portal Migrating Desktop. The RAS does the job submission and data handling on the Grid on behalf of the user. Several deployments of the web service RAS will be maintained by the partners involved in the task. In particular the RAS is able to use the workflow manager Kepler (maintained in this task by PSNC through the Polish JRU), which is the workflow tool selected by the Fusion modelling community, to submit workflow jobs.

Furthermore the RAS currently contains plugins allowing access to gLite and UNICORE based resources, which makes the combination of RAS + Kepler + (gLite / UNICORE plugin) a very useful tool for generic applications. The effort will be dedicated to maintain the RAS servers and to the upgrade of plugins according to the evolution of underlying middleware. One backup instance of RAS will be installed at the central services of the National Network provider, in Spain to increase reliability.

**GRIDWAY** is a metascheduler to launch jobs on Grid infrastructures, which also supports workflows. It has fewer capabilities than RAS because visualisation and interactivity are not provided, but it is used in many Fusion workflows because of its ease to use. Several instances of the metascheduler Gridway will be put in service, one of them in the central services of the National Network provider in Spain, RedIRIS, to increase reliability.

#### 1.5.5.4.5 TSA3.2.5 MPI

The need for focused MPI support within the infrastructure for non-HEP communities, notably CCMST, F, A&A, ES, LS and Solid State Physics communities was established during the EGEE-III project. Such an activity needs dedicated effort on the operations side to improve the quality of the infrastructure and effort from some of the critical user communities to actively provide feedback on the ongoing operational support for MPI.

CSIC maintains the *mpi-start* middleware component originally developed and deployed extensively during the Int.EU.Grid FP7 project. This product was crucial in enabling MPI on the EGEE infrastructure. TCD led the first MPI Working Group (MPI-WG) in EGEE-II and delivered middleware components which allowed Resource Centres to easily deploy MPI including an initial SAM-based test suite. Both CSIC and TCD have a well-established relationship in working together on the deployment of MPI-enabled middleware.

Due to its widespread interest in high-performance computation, the CCMST community actively works on using parallelisation tools and developing them further. Several laboratories in the community have contributed to these activities in EGEE and are committed to doing so within EGI-InSPIRE, leading to software optimised for both high-performance and high-throughput computing platforms. This makes the members of the CCMST community (some of which like UNIPG, ENEA, CNR/Democritos, INFN which are also important partners of IGI) ideal candidates to take care of further development of MPI on the grid. On behalf of the CCMST community, they will engage in using MPI-enabled applications on the production infrastructure to ensure that the MPI service is fully functional for the whole CCMST community. CCMST already offers support on MPI usage to its community on a smaller scale through the COMCHEM Virtual Organisation. Priority will be given for identifying applications currently running on dedicated supercomputers that should be ported to the grid, porting them, and registering them in the Applications database.

Fusion applications will be tested with special emphasis on assessing their scalability on Grid. The high number of legacy codes that exist in Fusion and plasma physics will in general pose a particular challenge to these tasks due to the architectural and software constraints that need to be taken into account. The Fusion community are reliant in the advanced predictive modelling capabilities provided by MPI applications in the area of magnetically confined fusion plasmas which are needed to understand and design the next-generation Fusion devices such as ITER (the International Thermonuclear Experimental Reactor). Computer science efforts have also been directed towards developing tools to facilitate the integrated plasma-edge model within the framework of an integrated Fusion framework (see EUFORIA project description in http://www.eu-euforia.eu) The wide range of physical scales that needs to be approached, from the core to the edge, makes it necessary to apply the techniques of parallel computing to integrate the full reactor geometry in the simulation. This application is an excellent driver for MPI support in Grids.

The partners involved in this task will:

a) Assume responsibility for the integration of MPI into the infrastructure by maintaining, testing and certifying the integration of MPI related software components coming from the external software providers into the infrastructure;
b) Maintaining a suite of monitoring tests (i.e. SAM tests) to validate that the deployed MPI components are working correctly at a particular site;
c) Establish an 'MPI Support Unit' in the EGI Helpdesk to ensure that support issues from users and sites will be effectively handled and any issues relating to the software are recorded for resolution by the relevant software provider and any resultant monitoring tickets are dealt with;
d) Disseminate the successes of this MPI support activity to the EGI community, in particular the relevant VRCs, and provide updated training material for national or VRC oriented training events.

General MPI support for all communities will be delivered through an operational activity (CSIC+TCD) and through engagement from the application communities (CCMST + FUSION).

#### 1.5.5.5 TSA3.3: Services for HEP

Building on the powerful generic infrastructure of the underlying grids that they use, the LHC experiments have developed important complementary services, particularly in the areas of data and workload management, as well as in support for analysis services. Such services, which extend the capabilities of the infrastructure by exploiting knowledge of the experiment's computing model, data placement policies and information in metadata repositories, allow these massive international communities to maximise the benefit of the grids that they use.

This task will focus on essential optimisations of the existing solutions as well as improvements that are mandated by the experience of initial long-term production data acquisition. As in the past, such "innovations" are expected to be of benefit to many communities and understanding how the advances made could move to the mainstream will be an important element of the work undertaken.

It is foreseen that the effort in this task will be integrated into the Grid Support group at CERN, focussing on workload management, data management, monitoring, service deployment and operation issues driven by the production needs of the supported VOs. Specific examples include the provision, deployment and operation of scalable solutions to the experiments' needs for data distribution, detector conditions, data distribution and access – such as the deployment of distributed services to provide the caching of database information for the reconstruction and analysis of physics data , automation of the management of experiment specific services at the various sites and the development and deployment of efficient monitoring tools that are essential to ease the tasks of shift operators.

This activity is of particular importance now as we enter the exploitation phase of the world's largest scientific machine – the LHC at CERN – and will allow us to capitalize on the investment made by the European Commission through its co-funding of three phases of the EGEE project. This has resulted in large-scale production use of world-class Grid-based solutions by many key communities and has established Europe's leadership in this area. In the short to medium term, it is expected that this will lead to significant advances in our basic understanding of the Universe, whereas, in the longer term, major spin-offs, both related to advances in science as well as in Information Technology, can be expected.

#### 1.5.5.6 TSA3.4: Services for LS

**Distributed Medical Data Management**

The large availability of tools for data management has opened up a portfolio of scientific applications on grid infrastructures. A new generation of grid applications in the fields of medical imaging and public health informatics is now under deployment with the potential for high scientific impact and visibility. The support of the high level services allowing the management of distributed data is of utmost importance to the growing adoption of e-infrastructures in life sciences.

CNRS will ensure sustainability and distribution of the Medical Images Management (MDM) tool. This work includes an MDM software update (in particular DPM-DICOM plugin update to preserve compatibility with the versions of MDM), software packaging and distribution (in particular, package updates to preserve compatibility with the future version of the gLite middleware), and MDM user support. Additionally, the MDM client will be enriched with the functionality required by further community use cases.

**Integration of molecular biology core resources on the biomed VO**

For the users of the Life Sciences Virtual Organization, access to the core molecular biology databases (UNIPROT, PDB, GenBank) is absolutely critical. Building upon the recommendations and developments of the EMBRACE network of excellence (DG-RESEARCH, FP6), the adoption of web services enables interoperability. However, further work is needed to allow life sciences grid users to query the core biomolecular databases at EBI in an easy and transparent way. A service will be set up to regularly install updated versions of the databases on storage elements of the Biomed VO. This work will contribute to the collaboration between EGI and ELIXIR, the distributed research infrastructure for molecular biology.

**Middleware for grids integration and deployment of bioinformatics and biomedical services and data**

On the one hand, bioinformaticians are often reluctant to use complex systems such as grids. On the other hand, they are facing a large number of analysis tools and data to process. The WISDOM Production Environment (WPE) first developed to handle large-scale drug-discovery experiments has been extended and shaped following the standards defined by the bioinformatics community within the EMBRACE project. The environment comes with a set of services that allow simple and efficient integration of tools and data on multiple grid middleware and computing facilities. The goal of this task is to extend and improve the WPE with new features to provide a sustainable platform for bioinformaticians and biologists through which they will get a transparent and unique access to their services and data on the multiple middlewares supported by EGI.

The WISDOM production environment will manage jobs continuously on all the supported infrastructures that will be able to fetch and compute any task from any services integrated into the platform. Users will interact with the system with a set of web services interfaces so that the WPE will provide a service-oriented abstraction layer to the underlying middlewares.

#### 1.5.5.7 TSA3.5: Services for A&A

Through MPI-based computational simulation, the A&A community produces huge amounts of data that needs to be analysed. For instance, cosmological simulations produce large datasets that cannot be easily investigated without using parallel procedures to extract information from the dataset – either visually or by performing statistical analysis (e.g. halo finders, halo extraction, mass function computation and so on). Projects such as the Sloan Digital Sky Survey, and next-generation sky surveys, such as the Low Frequency Array (LOFAR), the Square Kilometre Array (SKA) etc., are planning to collect large amounts of raw data eventually resulting in several hundreds of terabytes of data in tabular form.

The VisIVO Server tool is a command line application for intuitive data exploration through multidimensional views of data, both from simulations or from sky surveys, which can be easily installed on any server with a database repository and benefit from using grid-enabled HPC resources. This work will integrate the VisIVO Server so that it can use grid-enabled HPC resources. This integration will allow the exploration and visualization of large datasets, produced on the grid by user communities, as VisIVO already supports many non-proprietary data formats. VisIVO Web, a web portal for data analysis that mainly runs VisIVO Server on a local system, will also be supported. Some recent examples of the work produced using grid resources are available (http://www.oact.inaf.it/visivo/).

VisIVO can also be used by complex workflows during the production phase on Grid, to monitor generated data and to give feedback on the run itself. This is made possible by properly integrating typical applications on HPC systems and applications on the Grid. VisIVO allows users to manage many kind of input data formats, and could be used across many scientific domains.

#### 1.5.5.8 TSA3.6: Services for ES

Data are a key part of any ES application. Currently, the ES community is using several interfaces to access Data and Metadata outside of the EGEE infrastructure by using grid-enabled database interfaces such as OGSA-DAI, Spitfire, GReIC and AMGA. The data centres have also developed service tools for basic research activities such as searching, browsing and downloading these datasets, but these are not accessible from applications executed on the Grid. This task will also enable these tools to be accessed from the Grid. In collaboration with GENESI-DR (Ground European Network for Earth Science Interoperations - Digital Repositories) this task will maintain and evolve an interface in response to new requirements that will allow data in the GENESI-DR infrastructure to be accessed from EGI resources to enable future research activities by this HUC**.** The GENESI-DR provides one approach to scientists accessing satellite and earth observation data located in different data centers for data processing on the Grid. The GENESI-DR work is aligned with an international initiative GMES (Global Monitoring for the Environment and Security) that is supported jointly by the EC, the European Space Agency and their respective member states.

These data-access interfaces are critical for the VOs in ES, A&A, climate modelling. meteorology and projects from other domains (e.g. biodiversity projects such as LifeWatch, cyclops**,** civil protection, health and agriculture).

#### 1.5.5.9 Deliverables

| Deliverable Number | Deliverable Name | Work package Number | Nature | Dissemination | Delivery Date |
|---|---|---|---|---|---|
| DSA3.1.1-3 | Capabilities offered by the HUCs to other communities | SA3 | R | PU | M01,M13, M25 |
| DSA3.2.1-3 | Sustainability plans for the HUC activities | SA3 | R | PU | M09, M21, M33 |
| DSA3.4.1-3 | Annual Report on the HUC Tools and Services | SA3 | R | PU | M11, M23, M35 |

#### 1.5.5.10 Milestones

| Milestone number | Milestone name | Work package(s) involved | Expected Date | Verification |
|---|---|---|---|---|
| MSA3.1.1-3 | HUC Contact points and the support model | SA3 | M02, M14, M26 | Report establishing the HUC contact points and the support model for each of the tasks within the activity. The report will also detail how each task will gather and prioritise requirements from other users, and how they will be able to access the supported software/services. |
| MSA3.2.1-3 | Training and dissemination event for all shared services. Other activities within SA3 may wish to participate. | SA3 | M06, M21, M33, | Event advertising and other dissemination material |
| MSA3.3 | Hydra service deployment on a multi-servers configuration | SA3 | M12 | Service availability advertised |
| MSA3.4-1.3 | Medical Data Manager release | SA3 | M12, M24, M36 | Release made available for download |
| MSA3.5.1-6 | HUC Software Roadmap | SA3 | M03, M9, M15, M21, M27, M33 | Report describing the roadmap and release plans from the software development activities taking place within SA3. |
| MSA3.6 | Integration of the VisIVO server with the production infrastructure. | SA3 | M12 | Software release made available with supporting documentation and demonstration material. |

#### 1.5.5.11 Risk Assessment and Mitigation

| Risks | Impact | Probability of Occurrence | Mitigation |
|---|---|---|---|
| Services provided by one community on behalf of another are not responsive to other communities needs. | The other communities develop their own solutions. | Medium | Each shared service/tool will define a process by which new requests are recorded, assessed and prioritised. This will apply to both their own and other communities. Transparency in the decision making process will be assured by inviting members of the other user communities to be involved in these decision making processes. |
| The UMD services that the user community is building upon do not provide the coherent functional interface that the HUCs need. | Retaining the current users and attracting new users will be difficult if UMD is seen to have no coherence or integration. | Medium | The HUCs are represented on the MCB (See NA2) which will define the UMD Roadmap and assessment criteria. They will be able to also define their requirements as to the functionality that should be incorporated into the production infrastructure in the future, and select the software providers to deliver that software. |

| | | | |
|---|---|---|---|
| Communities develop and adopt different services that have similar functionality. | In the worse case the production infrastructure has to deploy different services with the same functionality, thereby multiplying the cost of supporting different user communities. | Medium | Within EGI-InSPIRE any duplication is understood. Within the broader community this could still happen. However, the open process within EGI provides more opportunity for communities to collaborate on services, and for EGI to make it clear the functionality that will need to be supported in new software for it to be deployed. |
| The production grid infrastructure does not meet the needs of the current HUCs and new potential HUCs (i.e. ESFRI projects). | The user base reduces as existing users go away and no new users are attracted to the infrastructure. | Medium | The HUCs will be represented in the management bodies guiding the development of the current production grid infrastructure. Technical outreach within SA3 and engagement within NA2 and NA3 will continue to show the benefits of the infrastructure for new communities. |