



Applications & Community Support: Transition Plan

1. Introduction

1. The usage of Grids continues to grow, both in terms of active communities as well as in resource utilization: the current European Grid infrastructure EGEE consists of more than 270 sites distributed in 54 countries. More than 140000 CPUs and 20PB of disk space serve the research purposes of more than 15 communities. Today there are more than 200 VOs and about 16000 active users;
2. The WLCG community, which has been running in continuous production mode for several years and is the leading user of the European Grid infrastructure, recently carried out a Scale Testing for the Experiment Programme'09 (STEP'09) exercise. As part of this programme the LHC experiments ran over 2 million analysis jobs and generated over 6 GB/s of sustained Grid traffic. The exercise was prepared with the experiments during a period of about two months and executed in two intensive weeks of 24x7 operations;
3. The EGEE Grid (as well as additional Grid resources used by WLCG, such as the Nordic DataGrid Facility in Europe and Open Science Grid in USA) has already achieved a high level of scalability and reliability in terms of services, software and operations, with service targets for sites being regular met or even exceeded (99% service availability being required for the most critical Tier0 and Tier1 services);
4. A stable and reliable infrastructure makes the Grid an attractive environment for both research and business activities and will help attract new communities: the Grid support team at CERN concurrently supports some 5 additional communities including simulation toolkits, United Nations Initiatives, theoretical physics studies, etc;
5. Based on this infrastructure, private enterprises, businesses and United Nations initiatives have shown their interest in Grid technologies since several years;
6. Profiting from the WLCG experience and know how, the Grid Infrastructure plays or will play an important role in Photon science etc. for the computing infrastructure of the ESFRI-Projects FAIR, EuroFEL and the XFEL with DESY as one of the leading partners.
7. After three successful phases since 2001, the EGEE project will finish in May 2010. The next Grid project EGI foresees significant changes to the current Grid environment that are described elsewhere in this document. The basic requirement of all communities is to ensure the stability of the current Grid infrastructure during and after the transition. This issue also affects the service level and operations procedures;
8. Europe's leadership not only of the development and deployment of grid systems but more importantly in terms of their successful utilization for



July 29 2009 – Draft 0.86

enabling science and other disciplines is well known and attested to by the numbers listed above. This leadership translates into a competitive advantage both for these communities as well as the clear potential for direct and indirect spin-off.

9. In this document we will address the most important aspects of a smooth transition from the point of view of HEP and related communities. The operations procedures and the required level of the services will be also detailed;

2. Areas of Key Impact

10. The expected results of this work are a marked increase in the number of Grid users, as usage expands from the data processing activities that have dominated until now into the realm of data analysis, scientific discovery and publication. This will be accompanied by wider inter-disciplinary collaboration, both through science (i.e. related disciplines) and technology (e.g. grid tools);
11. This can only be achieved by a significant simplification of user interaction with the Grid, through further adoption of existing tools such as those described in detail below, and by a flexible and scalable end-user support model;
12. This includes the establishment of community support, whereby the communities are encouraged and enabled to be largely self-supporting, with expert guidance to establish and optimize the support structures and associated tools. This is essential not only to deal with the large expansion in terms of number of users but also for long-term sustainability;
13. These items will have a positive and measurable impact on the quality and effectiveness of the e-Infrastructure which in turn will lead to corresponding benefits to the research communities that use it;
14. These activities will help to ensure Europe's leadership role in the areas of grid design, deployment and efficient exploitation.

3. HEP SSC: Motivation and Plans

15. High Energy Physics (HEP) community is the initiator and the leading user of the EGEE infrastructure (80% during the past year), as well other major Grids [references], including OSG and NDGF;
16. Along with production usage throughout the year, two specific programmes aimed at demonstrating readiness at the scale required for data taking, processing and analysis at the Large Hadron Collider (LHC) at CERN have taken Grid computing to new heights in terms of service availability, number of jobs per day, number of unique users, inter-site data transfer rates, as well as functionality and usability;
17. As a direct result of this work, a variety of other disciplines – linked either by the science that they perform, or else via the (Grid) technology that they use – have benefited measurably;



July 29 2009 – Draft 0.86

18. These activities have been supported during the past years through co-funding from a variety of sources, including national Grid projects and funding agencies, international projects supported through the EU and other agencies, as well as specific institutes such as CERN, which have allowed a relatively small “Grid support” team to empower a number of (very) large distributed communities;
19. To fully realize the value of this investment, this “Grid support” function must continue and it is foreseen that this be done via a “Specialised Support Centre” (SSC) for these communities;
20. We describe below the key functions and benefits of this SSC, together with a concrete transition plan and timeline;
21. The impact will include closer collaboration between scientific disciplines – HEP, astro-particle physics, fusion, light source or “photon” science – as well as numerous international communities, using user-friendly high-level interfaces to underlying Grid middleware services that simplify significantly the use of powerful e-Science infrastructures developed for the above domains;
22. Although the intent is to continue to support a variety of disciplines, we will refer to this activity as “the HEP SSC” in the remainder of this document.

4. WLCG Transition Plan

23. One of the key users of the services from the SSC that this section describes will clearly be the LHC experiments and the Worldwide LHC Computing Grid (WLCG) project;
24. Given the schedule for the restart of the LHC in late 2009, with a data taking run continuing more or less uninterrupted until late 2010, the transition must start prior to any FP7 and/or related funding being available;
25. This has been commenced through the standard CERN fellowship programme, through which two post-doctoral short-term positions have been opened as of summer 2009. However these positions only partially compensate for staff loss through end of contract / (external) funding and other departures. Further significant reductions will occur naturally as EGEE III draws to a close;
26. These reductions and indeed any further loss of effort represent a major risk not only to the communities supported but also to the use and adoption of grids in general. It is felt essential not only for these reasons but also to retain Europe’s leading role in e-Infrastructures that this concern be addressed with priority;
27. Notwithstanding the above, the existing effort is being refocused on the needs of the imminent startup of the LHC, including increased support for data analysis, where “community support” – whereby the large distributed community is trained and enabled to be largely self-supporting, backed up by a team of experts – is proving to be an attractive model;
28. For the purpose of this document we describe the proposed activities in terms of “work packages” that focus on specific communities, as well as “horizontal activities”, including training, dissemination and common tools;



July 29 2009 – Draft 0.86

29. Although expected to be (largely) performed outside the strict context of an SSC, we also enumerate our requirements in terms of middleware (release and service issues) as well as operations and user support;
30. Given WLCG's multi-national dimension, non-EGEE issues, including the interoperability and interoperation with NDGF and OSG, must also be addressed.

4.1. Operations and User Support

The WLCG requirements in terms of Operations and User Support are largely covered in the document “WLCG Operations for LHC Data Taking”. We reproduce here the conclusions:

31. WLCG Operations has reached stability and sustainability. For relatively low but non-zero cost it can be maintained and enhanced;
32. Manpower to support the key operations tools and to fill needed operations roles – both WLCG and experiment-specific – is required. It is expected that this be jointly funded to allow the successful exploitation of the world-class grid that has been built up over many years, together with international partners that make this a truly global enterprise;
33. We believe that WLCG operations experience and procedures can have significant value to other communities and are keen to share this knowledge.
34. Dedicated effort will ensure bidirectional communication between operations experts and the WLCG support team.

Manpower: 1 FTE liaison activity between EGI core operations and WLCG.

Call: 1.2.1.2.

Community: WLCG and International communities

4.2. Middleware Requirements

[This section – as with essentially all others – clearly needs to be synchronized with related transitions plans and proposals, e.g. EMI.]

The primary requirements are related to the release and updates of middleware versions and Grid service requirements. Given the move from deployment to production, the primary requirement is for middleware that is designed with robust service deployment in mind: this requirement must be taken into account from the early design stage and must be reflected in issues such as consistency and clarity of error messages and logging, the provision of the necessary hooks for monitoring and eventual debugging, as well as design for robust deployment (scalability and failover). In addition to these basic requirements, the schedule of large communities such as WLCG needs to be considered. In this specific case they are driven by that of the LHC machine, which will typically operate from Spring until Autumn annually (except at least during the startup years 2009/2010, when it is expected to run more or less continuously from late 2009 for about one year). During this period, major middleware or service changes cannot realistically be deployed in production (whereas minor updates and / or bug fixes can). In an ideal situation, larger changes would be made available such that they can be fully tested and debugged to allow production deployment several months prior to the annual accelerator startup.



July 29 2009 – Draft 0.86

Changes which miss this time window are unlikely to be deployed prior to the following shutdown period.

Although the schedule is expected to vary with discipline and associated scientific machine, similar constraints can be foreseen for other large communities.

Impact: It is crucial that operations and middleware teams work closely together. The middleware experts should provide the operations team with recommendations procedures to take the maximal benefit of the middleware. In addition the middleware has to be provided with the following rules in mind:

1. Experiment requirements in terms of new or improved features;
2. Time constraints per year to provide the new versions (see above).

Call: 1.2.1.2

Manpower: 1 FTE liaison activity. Manpower for 1.2.1.2 depends on services included in this item. Should a significant fraction of “higher level” services – such as FTS – need to be covered in this area the manpower requirements could easily rise to 10 FTEs or more. This is in addition to those that would be required for key experiment-specific services, such as those listed below.

Although this is an area still under active discussion, it looks likely that at least some fraction of the middleware components and associated services that have until now been co-funded through EGEE – e.g. FTS – will need to be covered in this area. We therefore use an estimate of 6 FTE (out of the possible 10 or more FTEs mentioned above).

Call: 1.2.1.2

Manpower: 6 FTE for essential common middleware & services for large existing user communities that are heavy users of DCIs and have multi-national dimension.

Community: WLCG and international communities

4.3. Service Requirements

The following services are expected to be required by all communities and are included here for completeness. The manpower requirements are estimated elsewhere in this document:

- GGUS: Relies on connections to local support ticket system – today in ROCs and sites
- Operations and Service Coordination: Currently CERN + ROCs
- ROCs: Support effort (TPM, COD)
- GOCDB: configuration DB, important for all configurations and definitions of sites and services
- Deployment support: m/w deployment/testing/rollout/support and PPS
- Operational Security coordination
- Policy development
- Accounting: APEL, infrastructure/DB and service and Portal (CESGA today)
- CIC: contact information, VO-ID cards, broadcast tool, automated reporting



July 29 2009 – Draft 0.86

Impact: The mentioned services are required for a functioning e-Infrastructure.

Community: All.

Call: 1.2.1.1

4.3.1. Additional Services for Large Existing Communities

In addition to the services listed above, WLCG also requires the following services. Most of these have been provided as part of the basic EGEE infrastructure and many developed and maintained with EGEE funds:

- FTS and SRM
- WMS
- Compute elements: CREAM-CE, ARC-CE
- Myproxy
- Remote UIs (WNs)
- LFC
- DB services and DB applications for both individual for the experiments use and also supporting other Grid services as FTS, LFC, etc.
- ENOC: Coordination of LHC OPN operations, currently by IN2P3
- VOBOXES and UI: Grid access services
- Availability/Reliability: SAM framework and migration to Nagios, SAM tests, Gridview and algorithmics, GridMap and MSG
- Dashboard and monitoring infrastructures: service, framework and common services, experiment-specifics
- Analysis Infrastructure
- Mass Storage solutions (Tapes, dCache, Castor, StoRM)
- Data access infrastructure (local access and high level libs. Also experiment specific solutions, a la xrootd)
- Information system infrastructure (including BDII) and associated query tools
- Security infrastructure (also to the level of jobs, a la glxexec/SCAS)
- VOMS

In addition the specific services required by the experiments are:

ATLAS

:

- The Ganga framework, developed in common with LHCb, is used to submit and monitor distributed analysis jobs on the Grid
- The PanDA production and job distribution system
- The DDM (distributed data management) system used by ATLAS

(CMS' requirements are quantitatively similar but different in detail due to fundamental differences in the computing models of these two large collaborations, e.g.

- ProAgent for production
- CRAB for analysis
- Phedex for transfer.)

In the transition from EGEE to EGI and the related SSC infrastructure, ATLAS expects that generic packages mentioned in section 4.3 will be supported, maintained



July 29 2009 – Draft 0.86

and evolved in the context of the Middleware unit whereas specific services for the experiment and for large communities will be supported by HEP dedicated funding (in call 1.2.1.2 "service deployment" and in the SSC call).

In addition to the effort for the development and maintenance of generic and specific middleware packages, currently EGEE funds effort for VO operation support. The GS group at CERN is essential for the efficient and smooth operation of the global system and to provide user support.

ATLAS effort in EGEE: ATLAS is currently supported by 8 people, directly or indirectly funded by EGEE at CERN (5 staff, 2 fellows and 1 technical student), who work on the development of middleware layers and operations support. Their work is essential to ATLAS. Their number has already decreased recently as the end of the EGEE-III project is getting close, but we expect the number of these people not to decrease further during the transition phase and the new EGI/SSC funding scheme.

In terms of ALICE (again, the requirements of LHCb are quantitatively similar), this experiment uses general middleware services for workload and data management, several storage solutions and ALICE-specific top level services (denoted as AliEn) as user interface to Grid, distributed catalogue, job queues and prioritization. The list of services is given below with potential call assignment:

Type of Service	Name	Current provider	EGI Call	Remarks
Workload	gLite WMS	gLite	1.2.1.1	
Workload	CREAM-CE	gLite	1.2.1.1	
Workload	LCG-CE	gLite	1.2.1.1	Deprecated, used only on sites with no CREAM-CE
Data	FTS	gLite	1.2.1.2	
Data	CASTOR2 dCache, DPM with xrootd	General Storage solutions	1.2.1.1	
Security	VOMS, VOMRS		1.2.1.1	
Information	BDII	gLite	1.2.1.1	
Ticket system	GGUS		1.2.1.1	
UI	gLite-VOBOX	gLite	1.2.1.2	
Data	SRM		1.2.1.1	
Monitoring	SAM, Dashboard	WLCG	1.2.1.2	
Experiment framework	AliEn	ALICE		



July 29 2009 – Draft 0.86

ALICE requests continuous VO Grid support, as already provided by the CERN IT/GS and IT/DM teams in the EGEE/WLCG. In terms of manpower this is a 2 FTE equivalent, which should be included in the manpower call 1.2.3.

Community: WLCG

Call: 1.2.1.2

Manpower for 1.2.1.2 depends on services included in this item. On the assumption that this includes the main services that are currently run on top of the basic middleware services – which are expected to be covered in 1.2.1.1, but which may also need to be covered in 1.2.1.2 – ~10 (additional) FTEs for experiment-specific services (building on those provided by the 6 FTEs described above).

A good example of the need for experiment-specific services came from the recent WLCG STEP'09 Post-Mortem workshop. Whilst the LHC experiments are using FTS services – established at the WLCG Tier0 and all Tier1 sites – these services are necessarily generic and cannot know about an experiment's computing model, specific topologies or indeed the data itself. Experiment-specific solutions, however, whilst building on this highly reliable and powerful infrastructure, can make certain optimisations which take into account this additional knowledge. For example, PhEDEx, the CMS data movement system, is able to source files belonging to a larger dataset from alternative sites, leading to additional robustness and performance. As much as 50% of the data – possibly more – may be retrieved from an alternative source: functionality that cannot be provided at the FTS layer in a non-intrusive fashion.

5. Use case 1: High Energy Physics (HEP) – WLCG

HEP and in particular WLCG is the largest community of the EGEE infrastructure, both in terms of number of users and resource utilisation. WLCG also has correspondingly large requirements in the number of concurrent and completed jobs per day, I/O and bulk data transfer rates, custodial storage and so forth.

It is the community that often performs (pre-) production testing of new services and / or releases – benefiting not only WLCG (and HEP) but also many other Grid communities.

The complex computing models of the LHC experiments have implications on the operations procedures and Grid support infrastructures. The policies and procedures established for the HEP communities – e.g. within WLCG – have frequently been adopted as part of the standard EGEE procedures.

WLCG can therefore be considered as a reference for other scientific communities. Collaborative activities between WLCG and other communities have been supported during the past years through co-funding from a variety of resources, including national Grid projects and funding agencies, international projects supported through the EU and other agencies, as well as specific institutes such as CERN, which have allowed a relatively small “Grid support” team to empower a number of (very) large distributed communities.

A smooth transition from EGEE to EGI is essential for the WLCG community. A transition that satisfies this community can also be successfully applied to others. Given the schedule for the restart of the LHC in late 2009, with the data taking run



July 29 2009 – Draft 0.86

continuing more or less uninterrupted until late 2010, the transition must start prior to any FP7 and/or related funding being available.

In terms of WLCG application support, there is a well established and trusted infrastructure that should be kept largely as it is during the transition phase to the SSC infrastructure (although in terms of manpower it is already going down as people on short term contracts find employment elsewhere). A specialised support team called EIS (Experiment and Infrastructure Support) placed at CERN works daily with the experiments in order to ensure a good implementation of the Grid infrastructure into their computing models. The migration of the functions of this team to the future SSC as the support infrastructure of the WLCG experiments is crucial and supported by the communities. Currently EIS is focused on gLite middleware support, offering only limited assistance to WLCG sites that use ARC middleware – this can be improved in future. The principal roles of this team are the following:

- Liaison between the WLCG community and services, operations, middleware teams and the sites;
- EIS provides Grid expertise for the middleware implementation into the application computing model. This might include the development of specific application software towards the final implementation and the creation of additional components which might later be shared by other communities (as has happened in the past). In addition this team is responsible for the evaluation of new middleware functionality and to perform scalability tests based on the needs of the community;
- The EIS team manages also VO-specific services and provides user-support as required (e.g. submitting and following up on GGUS tickets on behalf of the VO);
- This team has the expertise to solve Grid related issues for running applications both to the site and middleware level. The members have a deep knowledge of both Grid infrastructure and WLCG computing models;
- They are responsible for application-oriented documentation and training;
- Using this knowledge they also create high level Grid implementation tools;
- EIS is responsible for experiment-specific tests executed within the overall monitoring framework (SAM);
- The team also provides links to operations activities, such as creation and registration of new VOs, VO manager activities, negotiation of resources and services at the sites, etc.

Impact

This will allow the WLCG community to successfully exploit the scientific potential of the world's largest scientific machine, potentially leading to new levels of understanding of the Universe around us, in particular in areas not fully or adequately explained by the so-called "Standard Model".

Effort

8 FTE support personnel for the 4 main LHC experiments.

Technical consultants for liaison with operations and middleware (2 x 1 FTE – already counted above)

Expected VO-specific operations effort provided by the T0, T1 and other sites to cover the LHC experiments' needs (estimated at 1 FTE per VO per site).



Deliverables:

PM12: to be agreed with community

PM24: as above

PM36: as above

Community: WLCG and international communities

Call: 1.2.3

5.1. Analysis Support for HEP

From Q3 2009, the user analysis tasks in WLCG will become especially important and are key to achieving the final goal of the experiments. All LHC experiments have a user analysis infrastructure with the following characteristics:

- Analysis jobs will run where the data are physically stored;
- The user analysis frameworks typically use the same infrastructure used for data re-processing;
- User analysis requires transparent access to all resources (independent of the middleware stack(s) involved) and / or data location.

For three of the LHC experiments, ATLAS, CMS and LHCb, the user analysis frameworks have been developed and supported by the support infrastructure of HEP in EGEE. These frameworks are Ganga for ATLAS and LHCb and CRAB for CMS. In addition the Dirac user analysis infrastructure for LHCb and the AliEn infrastructure of the ALICE experiments are well known and supported by members of the EIS team. In particular in the case of the ALICE experiment, the data access infrastructure has been partially developed and maintained also within the EGEE-III project.

All the frameworks listed above incorporate support for several middlewares (gLite, ARC, OSG) and batch backend setups, and this complexity constitutes one of the toughest challenges for developers and supporters. While current EIS expertise is concentrated in gLite, knowledge of ARC and possibly other middleware aspects is essential for many HEP community members. Due to this community nature, there always will be power users working outside the tools for which providing good support will have extra benefits.

The following tasks should be included:

- Development and maintenance of the user analysis frameworks
 - Ganga, CRAB
 - Monitoring tools associated to the analysis jobs: Dashboard
- Support and services in terms of data access
 - Xrootd
 - High level tools
- Grid support infrastructure (EIS)
- Middleware expertise and liaison between experiments and middleware providers
- Training actions
- User support



Impact

Enable several thousand physicists worldwide to process and analysis data using the grid.

Effort: 2 FTEs

Deliverables:

PM12: to be agreed with community

PM24: as above

PM36: as above

5.2. Dashboard

The Experiment Dashboard [1] monitoring system was developed in the framework of the EGEE NA4/HEP activity. The goal of the project is to provide transparent monitoring of the computing activities of the LHC VOs across several middleware platforms: gLite, OSG, ARC.

Status:

Currently the Experiment Dashboard covers the full range of the LHC computing activities: job processing, data transfer and site commissioning. It is used by all 4 LHC experiments, in particular by two largest, namely ATLAS and CMS. Generic functionality, such as job monitoring, is provided by Dashboard server to all VOs which submit jobs via gLite WMS.

The Experiment Dashboard provides monitoring to various categories of users:

- Computing teams of the LHC VOs
- VO and WLCG management
- Site administrators and VO support at the sites
- Physicists running their analysis tasks on the EGEE infrastructure

To allow for the development and building of components of Dashboard Monitoring applications, a Dashboard framework was designed. This is currently used by other projects and development teams and not exclusively for the development of the monitoring tools. The Dashboard framework is used as well for the construction of the high level monitoring system which provides a global view of LHC computing activities across all LHC experiments, both at the level of the distributed infrastructure in general as well as on the scope of a single site.

Web monitoring shows heavy use of the Dashboard servers, for example the dashboard of the CMS VO serves 2300-2500 unique visitors per month with about 30K pages accessed daily. These numbers are growing steadily.

Future Evolution:

The future evolution of the project is driven by the requirements of the LHC community which is preparing for LHC data taking at the end of 2009.

The main strategy is to concentrate effort on common applications which are shared by multiple LHC VOs but can also be used outside the LHC and HEP scope.

Impact:



July 29 2009 – Draft 0.86

Reliable monitoring is a necessary condition for production quality of the distributed infrastructure. Monitoring of the computing activities of the main communities using this infrastructure in addition provides the best estimation of its reliability and performance.

The importance of flexible monitoring tools focusing on the 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 like large simulation campaigns in HEP or drug searches in BioMed) 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 service provided by the infrastructure. Single users are typically scientists using the Grid for analysis data, verifying hypothesis on data sets they could not have available on other computing platform. In this case the monitor is a guide to understand the progress of their activity, identify and solve problems connected to their application.

This 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).

Effort

4 FTEs

Deliverables

PM12: Status report

PM24: Status report

PM36: Status report

Community: WLCG, international communities and rest of SSC

Call: 1.2.1.2

5.3. Ganga

Ganga is a flexible and user-friendly tool to enable large user communities to move away from a niche usage of the Grid. It supports the access to major grid infrastructures (notably EGEE, OSG and NDGF) and to all major batch systems. EGEE has supported the Ganga project via the EGEE NA4 HEP activity.

Detailed information on the Ganga system has been described in a recent publication [2]. Users can seamlessly run their applications by combining the usage of both Grid and batch resources. It is a genuine tool to run applications on multiple Grids (interoperability) and to protect users from changes of the middleware layer (new APIs). Ganga’s strengths are in its non-intrusive nature, allowing its users to run a applications without change on the Grid.

Status:

The core of Ganga is mature and which allows the development activity to be focussed on experiment specific extensions and plug-ins for ATLAS and LHCb.

Impact:



July 29 2009 – Draft 0.86

Ganga is one of the most useful elements in our toolbox to support large user communities, enable new ones and collaborate with other infrastructure projects. The return of investment is ensured by the large user communities, by the existing and future activities in the area of support and by the outreach to new communities.

Effort:

Engineer: 2 FTEs

Deliverables:

PM12: Status report

PM24: Status report

PM36: Final report

Community: WLCG, International communities and rest of SSC

Call: 1.2.1.2

5.4. DIANE

DIANE is a tool supported by NA4/EGEE project and listed in the RESPECT program [3].

The tool improves the efficiency of using distributed resources, especially Grid, by providing private, user- or application-specific overlays for handling workloads in Master/Worker paradigm (aka pilots, agents, glide-ins). DIANE may be used to easily create robust production-like systems running in an autonomously over long periods of time

DIANE is used in conjunction with Ganga [6] to provide an advanced grid-access and workload-management solution for smaller experiments and user communities with zero deployment and zero customization of the infrastructure required, allowing a seamless combination of Grid and local resources. The DIANE/Ganga system is used directly by individuals in several application domains such as medical image analysis, MRI cardio-vascular analysis, astrophysics simulations, fusion reactor applications, high-energy physics experiments (COMPASS and Atlas for Athena-based analysis on local clusters).

DIANE/Ganga system is used as a building block for community-specific services such as Grid portals, as demonstrated with ASGC bioinformatics Grid Portal for genomics and drug design. The combination of both tools also provides an attractive R&D platform (easily pluggable schedulers and parameterization of workload execution) which attracts researchers in various communities, including middleware solutions (e.g. SAGA) and application models (workflow management with DAGMan and Vbrowser/Moteur)

Status:

The project is mainly in maintenance mode: bug fixes, porting to emerging platforms, integration of components contributed by user communities, occasional new features and extensions, updates of documentation etc.

Effort



July 29 2009 – Draft 0.86

An effort (approx. 1FTE) is required to maintain the system in the current state and to provide expertise on application porting.

Impact: Diane is serving a long amount of communities which come to the Grid to use their resources in a very scheduled period of time. All those communities using the Grid in a scheduled way (as most of the high CPU demanding communities as most of the United Nations Initiatives) could eventually not fulfilled their requirements in terms of required CPU time if a resource optimization infrastructure in not included in the project.

Community: international communities and rest of SSC

Call: 1.2.1.2

5.5. AMGA

AMGA [7] is a general purpose metadata catalogue and part of the gLite Middleware distribution. It was developed within the HEP support activity in the EGEE project, when it became clear that many Grid applications needed metadata information on files and to organize a work-flow [8]. AMGA has been developed and supported by a collaboration of CERN, INFN in Catania and KISTI in Korea and is currently supported by INFN and KISTI.

AMGA as a metadata services allows users to attach metadata information to files stored on the Grid, where metadata can be any relationally organized data typically stored in a relational database system. In addition, the metadata in AMGA can also be stored independently of any associated files, which allows AMGA to be used as a general access tool to relational databases on the Grid. AMGA features a simple to learn metadata access language in addition to SQL, which has been very useful for the adoption of AMGA in smaller Grid applications, as it lowers the technical hurdle to make use of relational data considerably.

One of the main features of AMGA and unique to it, is the possibility to replicate metadata between different AMGA instances [9] allowing the federation of metadata (e.g. by the Health-e-child¹ project), but also to increase the scalability and improve the access times on a globally deployed Grid (as done by the Wisdom project). In addition to performance, security was another major concern during the design and development of AMGA [10] and AMGA features different authentication and access control methods. These advanced security features have made AMGA the de-factor standard for metadata and relational database access on the Grid for biomedical applications (e.g. [11]), within and outside of the EGEE project. Prominent projects making use of AMGA in this field are the Wisdom², and the Health-e-Child projects.

The easy usage and introduction of AMGA in Grid-enabled projects have made AMGA important components in many efforts to Grid-enable applications, like the Unosat project³ or the SRM-SRB⁴ Grid storage project used for digital libraries.

¹ An integrated platform for European paediatrics based on a Grid-enabled network of leading clinical centres, <http://www.health-e-child.org/>

² An initiative for grid-enabled drug discovery against neglected and emergent diseases <http://wisdom.eu-egge.fr/>

³ A UN-CERN collaboration to give field-workers in disaster areas access to satellite imagery and their treatment on the Grid: <http://unosat.web.cern.ch/unosat/>



July 29 2009 – Draft 0.86

In spring 2009, AMGA featured also the first implementation of the proposed WS-DAIR standard proposed by the OGF for relational data access on the Grid. A thriving collaboration between the EGEE community and the OGF community entailed, aimed at verifying the standards proposal and evolving it towards future needs.

Status:

The project is effectively in maintenance mode, where the effort for fixing bugs and code maintenance is provided by KISTI as part of their contribution to the gLite Consortium.

Impact:

In addition to the code maintenance it will be necessary to provide support for the existing user communities and continue to in supporting the adaptation of new applications to the Grid, both SuperB and the SuperBelle B-Factory experiments have voiced their interest in AMGA. In addition, the requirements of these users need to be represented in the standardization efforts. In addition we advocate the importance to devote a small yet visible effort which should be present as back-office effort for allowing the usage of AMGA in the above-mentioned activities (training, interoperability and gridification).

Effort:

Engineer: 1 FTE

Deliverables:

PM12: Status report

PM24: Status report

PM36: Final report

Community: projects besides WLCG as biomedical communities and international applications.

Call: 1.2.1.2

5.6. Grid technology outside LHC

Grid technology is seen as a major compute paradigm for the LHC data analysis. Nevertheless Grid technology is heavily in use in other HEP experiments as well like in the data analysis of the HERA experiments, in the detector design for ILC (the ILC community relies on the (WLCG) Grid for their detector studies in the context of the EU project EUDET.), the theory and in astroparticle physics like Icecube. For all those experiments, the major centres like DESY&INFN are hosting a full set of Grid Services based upon EGEE middleware components. These are used by the various communities as follows: MC and data analysis for H1, HERMES, ZEUS; the CALICE collaboration uses the Grid explicitly to store their testbeam data, being regularly taken at DESY, CERN, and Fermilab.

Therefore the role of major HEP-centres like DESY&INFN in a HEP SSC is, to

1. design and operate a basic grid infrastructure for non LHC communities from HEP;

⁴ <http://www.beliefproject.org/news/srm-srb-interface-a-new-milestone-for-grid-interoperation>



July 29 2009 – Draft 0.86

2. to further develop the grid enabled storage element dCache;
3. to support the scientific users in deploying the grid technology.

Effort

DESY requires 2 FTE funded by the EU and will add 2 FTE from the own staff.

Community: HEP

Call: 1.2.3

5.7. Small and medium sized experiments at CERN

Outside the mainstream LHC experiments a number of small and medium sized experiments (SME) will either be already taking data or will start taking data within the time frame of EGI. Around 500 scientists in total are participating in these experiments, which cover a large spectrum of HEP. Included in this list are experiments like COMPASS, the rare Kaon experiment NA62 and the heavy ion experiment NA61. These more traditional HEP experiments will have needs similar to those of LHC experiments, with large Monte-Carlo productions, relatively large data samples of up to a PB and will be mainly data driven. On the other hand, experiments like AEGIS (Gravity interferometry on Antihydrogen), ASACUSA (spectroscopy of antiprotonic helium) and CLOUD (Cosmic ray studies on clouds) are interested in CPU intensive simulations of their apparatus and atomic processes for optimisation purposes. They are expected to have small IO needs, but would like to explore a large phase space of parameters to tune their setup.

Effort

2 FTE

Community: HEP

Call: 1.2.3

6. Collaboration with International Entities

In EGEE we observed that the best way to promote Grid computing was to identify established communities with immediate requirements. Building collaboration by successfully attacking a concrete problem has been shown to be an excellent way to get visibility and to promote broader usage of Grid technology and of the EGEE infrastructure in particular. Target communities of interest include those already operating at a large scale but not yet using large distributed computing infrastructures as a permanent component of their normal operational model, such as large scientific or technical communities and / or enterprises. The interest is that one can effectively reach large scientific and technical communities.

In the recent past CERN established a collaboration with ITU (International Telecommunication Union) during the 2006 Radio Communication Conference. The Grid was used to complement the ITU computing system to perform heavy optimisation studies (200,000 jobs per optimisation round to be executed in less than 12 hours). The results became eventually part of an international treaty. The usage of the Grid was presented during the same conference in a plenary meeting (185 participating countries, delegations with technical and diplomatic staff).



July 29 2009 – Draft 0.86

More recently CERN helped the launch of an FP7 project lead by UNEP (EnviroGrids: Black sea catchment studies – 27 partners most of them from the Black Sea coastline countries). The contribution of CERN (Grid support) is relatively small (X% of the project budget) but makes possible an essential part of the project's programme of work (systematic study of the hydrological model and data comparison). This is a quite effective way to ramp up such an ambitious project, allowing the project management to concentrate in building the collaboration and the participating scientists to share their diverse competences to face the project scientific challenges.

The goal of formalising the existing activities is to improve and stabilize contacts with international organisations such as WMO, WHO, ITU. Of particular interest would be to establish contacts with EU agencies like the European Centre for Disease Prevention and Control (ECDC), the European Chemicals Agency (ECHA), the European Environment Agency (EEA) and possibly others.

These large entities are essential gateways to the corresponding scientific communities. Several of these institutions have privileged contacts with strategic economical players (e.g. ITU and the world of the Telecoms, WHO and the pharmaceutical industry, etc...). In this fact resides the strategic interest of this activity. In addition these organisations can reach communities (especially small or developing countries) which are often otherwise excluded from innovation by digital and social divides.

This activity should be lead by CERN because of its established role in attracting and collaborating with other activities (scientific and technical) in Grid technologies. Equally important is the experience matured by CERN as centre of excellence in computing for the HEP community (hence one of the largest and probably the most organised international research community).

Effort:

Activity leadership (CERN): 50% of a FTE

Technical consultant (CERN) 100% of a FTE

50 kEuro travel budget

50 kEuro the organisation of events

Deliverables:

PM6: Establish a cross-fertilisation plan with one of the large partners (to be used as a template for the others)

PM24: Report on the collaboration with at least 2 partners

PM36: Final report of the collaboration including prototype activity and a plan to maintain the established collaborations

Impact: The collaboration of WLCG with other communities has been a strong point of the project in the last years. The continuation of these collaborations will attract many users to the Grid coming from United Nations agencies and serving therefore directly or indirectly to thousands of persons all over the world. In addition the continuation of this collaboration could lead to a future Grid United Nations platform which will host applications and activities of these agencies leaded and supported by the WLCG project.



Community: International communities

Call: 1.2.3

7. Use case 3: Photon Science

The community of users in photon sciences has many similarities to HEP-users which makes it worth to use the synergies but there are also some major differences. The user community itself is much more heterogeneous and thus the applications and the basic ICT-knowledge vary a lot. This requires more effort in training, operation and administration of grid resources in comparison to the HEP community. Furthermore the typical length of an experiment last in contrast to the large HEP-Experiments only day's or weeks and the experiment groups are much smaller. This leads to a different operation model which can be seen in a much more "burst mode" fashion. Some of the experiments do have large data volumes, which occur during an experiment. Those data have to be made accessible through grid methods, where again in contrast to HEP security issues can play an important role.

In the photon community there are already many ongoing grid activities, in particular at large Labs like DESY and in the EGEE context. With the ESFRI-Projects EuroFEL, XFEL – both at DESY and the ESRF-upgrade, major new facilities are set up, which will give a boost to this research field. Therefore a distributed computing infrastructure is needed, which rely on the tremendous experience and development, achieved in EGEE and in the HEP community. Nevertheless this has been adapted to the needs of photon science. In particular, the proposed tasks are in continuation and broadening of the current EGEE activities:

- 1) User administration and access through Grid methods
- 2) Large grid enabled Data storage
- 3) Data access and distribution through Grid methods
- 4) Operation of Grid resources with common EGI middleware
- 5) Training and education of Users
- 6) Analysis in Grid environments

Effort

The photon community will propose a SSC, which will be driven by 4 major "light sources" in Europe, which are as well involved in the relevant ESFRI-Projects.

These Centres will be:

- DESY with the ESFRI-Projects EuroFel and XFEL
- PSI also with the ESFRI-Projects EuroFel
- Elettra (tbc)
- ESRF with the ESFRI-Project ESRF-upgrade (tbc)

It is expected to have 2 persons at each site, one unfunded and one funded by this project.

For the placement of this effort and according to the proposed scope of the project, each partner has the following role:

- DESY: grid operation, storage elements, Analysis
- PSI: Operations, user interface, storage elements
- Elettra: middleware, training, operation



- ESRF: operations, user interface, training

Impact (tbc)

Deliverables (tbc)

8. Use Case 4: FAIR

The particle accelerator complex FAIR (**F**acility for **A**ntiproton and **I**on **R**esearch) in Darmstadt, Germany is one of the largest projects of the ESFRI Road Map. The FAIR Baseline Technical Report⁵ describing the accelerator complex as well as the experiments was authored by more than 2500 scientists from roughly 250 research institutions from 44 countries. 3000 scientists are expected to carry out experiments at FAIR each year. After multi-annual planning and preparation civil construction is expected to start in 2010. The first beam is expected in 2015/16. FAIR will serve about 20 scientific collaborations from four major fields of research and applications.

The four scientific pillars of FAIR are:

APPA: Atomic physics and applied sciences in the bio, medical, plasma ESA, and material sciences;

CBM: Physics of hadrons and quarks in compressed nuclear matter and antimatter;

NUSTAR: Structure of nuclei, physics of reactions, nuclear astrophysics and rare isotope beams;

PANDA: Hadron physics, antiproton physics, charm and hyper matter.

The computing and storage requirements for FAIR are expected to be about an order of magnitude higher than the requirements of the LHC experiments. Because of the later start, the overall complexity of the system may be lower due to advantages from Moore's law. An e-infrastructure evolving around a combined tier0/tier1 facility collocated at GSI and at the University of Frankfurt (about 30km to GSI) and integrated in the European grid infrastructure is planned to support the experiments.

CBM and PANDA will use FAIR in an HEP-like mode – huge detectors run by a single collaboration all over the beam period. But the data processing model will move away from the hierarchical trigger systems used at LHC. The experiments require very complex algorithms for event selection, not allowing for definitions of data subsets to be processed e.g. by a first level trigger only. Therefore they will transport the entire data stream from the event building network of the detectors into a processor farm. The two other communities will have a lot of smaller collaborations. The communities involved in FAIR are therefore much more diverse than the user communities from other large-scale research infrastructures. This unique sociology will be challenging for the efficient use of ICT infrastructure and requires transversal tools across all communities.

In the long run FAIR has to become an e-infrastructure. On shorter terms the construction of FAIR must be accompanied by a raised level of e-infrastructure awareness and usage. It is important to closely cooperate with WLCG. PANDA and CBM have already started using the grid for detector simulations. By the end of this

⁵ FAIR Baseline Technical Report: accessible via www.gsi.de/fair/reports/btr.html



July 29 2009 – Draft 0.86

project all FAIR communities will have developed and elaborated their ICT road maps.

Effort: 4 FTEs (1 per FAIR community)

Deliverables: tbd


Community: FAIR

Call: 1.2.3.

9. Common Tools

The HEP support team has invested significant effort in the creation of tools that ease access to and the management of Grid infrastructures. These tools have been adopted by a large number of communities due their high level of abstraction and the stability. Access to the Grid, optimal resource usage, monitoring of jobs etc. are managed by this set of tools for which continued support is clearly required.

As today we observe several hundred Grid users using Ganga in their daily activity, a number in steady rise over the last years. Interestingly enough several users (not belonging to the HEP community, sometimes even from outside EGEE) “discovered” Ganga and started using them (often in a complete autonomous way using the material available on the web) as shown in Fig.xxx. Note that out of the ~500 users



per month about 25% of the users are non-HEP users.

Fig.x: Time evolution of Ganga users per month. Every month, unique users from ATLAS (blue), LHCb (green) are counted alongside with users from other virtual organisation (light brown). The dip in Nov 2008 is due to a monitor failure.

Ganga is the main tool behind a large and very successful activity of application porting (“gridification”) performed by CERN. More than 20 applications have been ported ranging from Fusion, Material Sciences, Accelerator Studies, Biomedical applications such as dosimetry in oncology and drug searches, large-scale software testing and many others.



10. Training and Dissemination

Training and dissemination are activities that will continue in the context of the current related activities. In this document we mention some examples of training events performed and managed by HEP members

Training:

Over last few years an ecosystem of user communities and applications have appeared around Ganga and DIANE tools, and the dissemination and training have played an increasing role in our activities. The long-term vision for training is to off-load the developers and core team at CERN by creating a situation in which user communities may provide a self-sustaining user-support and take, partially, a role of the driving force for the tools. We receive notoriously good feedback for the quality of our tutorials, which provide not only practical introduction to the tools, but also inspiration and more insight to the students than other, often more formal training proposed in EGEE. Here is a list of major training events (outside of LHC context) in last 12 months:

- *Train the trainer*, CERN, 2008
- *How to use the Grid for physics and medical applications*, IEEE NSS 2008, Dresden, Germany
- *EGEE tutorial: Grid introduction*, CSC Helsinki, Finland, 2009 (invited)
- *Application Grid Enabling*, Baltic Grid All Hands Meeting, Riga, Latvia, 2009 (invited)
- *Best practices in putting application on the grid*, Baltic Grid Summer School, ITPA VU Astronomical Observatory, Lithuania, 2009

The tutorials are organized via different channels, including the contact to the EGEE training group, which also picked up (partially) our training program for generic EGEE tutorials organized by them.

The key of the success is that the training activity involves Grid experts with day-to-day experience in supporting user communities and tools. The expert involvement may be direct or via Train-the-Trainer concept such as the case of BG Summer School tutorial.

Dissemination:

Through EGI actions as:

User Forums

- International events (at least once per year) to bring together Grid users
- Exchange of experiences and techniques
- Scientific program coordinated by EGI which also need to ensure a high qualified technical presence to confront problems from different communities

Topical Workgroups

- Users facing similar problems while using the Grid infrastructure
- Meeting prepared to devise solutions

Discipline Meetings



July 29 2009 – Draft 0.86

- Associated to specific disciplines to discuss common solutions related to the use of middleware, techniques, data resources.

11. Liaison with other projects / workpackages

11.1. Summary of support team tasks

- Liaison between the community and services-operations and middleware teams
 - The application might not know many Grid related issues
- Provides Grid expertise for the middleware implementation into the application computing model
 - Software development towards the final implementation
 - Creation of missing components which might be shared by other communities
 - Evaluation of the middleware functionality and scalability tests based on the community needs
- Management of VO-specific services
- User-support if needed (a la GGUS)
- Expertise to solve Grid related issues for running applications
 - To the site and middleware level
 - Deep knowledge of both Grid infrastructure and VO computing model
- Documentation and training (application oriented)
- Liaison with the sites
- Creation of Grid implementation tools
- Creation of the community monitoring systems
- Operations activities if needed
 - Creation and registration of the new VO, eventually VO manager, negotiation of resources and services, etc

11.2. Inter-disciplinary Scientific Collaboration

In the context of a multi community HEP SSC that will include also photon science, astro-particle physics and Fusion applications, the current collaboration which has been created with these large communities in the context of separated clusters in EGEE will be highly promoted. Based on the experiences gained during the lifetime of the EGEE project, the unification of efforts for those communities which can share similar requirements or goals is a fundamental task which will need to be highly promoted in EGI. This will follow to clear benefits from the point of view of the man



power effort, middleware creation (middleware duplication should be highly discouraged) and definition of operation procedures.

In spite of the relevance of this issue, it should be treated to the level of SSC definition and setup rather than at the level of the transition phase. In this phase, although the basis of the future SSCs will be defined, the guarantee of a stable and consistence Grid environment will be the major issue that we will have to face.

11.3. Summary of Manpower Requirements

Activity	Effort 1.2.1.2	Effort 1.2.3
Operations and user support liaison	1	
Middleware liaison	1	
Services for heavy users (“core”, i.e. FTS)	6	
Additional VO Services	10	N/A
Dashboard, Ganga, Diane, AMGA	4 + 2 + 1 + 1	
WLCG “EIS”		8
Analysis tasks for WLCG		2
Grid technology outside LHC		4+2+4 + 4 - 8
International communities		1.5
TOTAL FTE	26 (+4?)	~21.5 – 25.5
TOTAL EU Funded	13 (+2?)	11 – 13

12. References

[1] “Dashboard for the LHC experiments”, J. Andreeva et al., CERN-IT-NOTE-2007-048, November 2007, in the Proceedings of Conference on Computing in High Energy and Nuclear Physics (CHEP07), Victoria, Canada, September 2007.



July 29 2009 – Draft 0.86

- [2] “Ganga: a tool for computational-task management and easy access to Grid resources”, J. T. Moscicki et al., accepted for publication (May 2009).
- [3] <http://technical.eu-egge.org/index.php?id=290>
- [4] J.T.Moscicki et al.: *Dependable Distributed Computing for the International Telecommunication Union Regional Radio Conference RRC06*, <http://arxiv.org/abs/0906.2143>, submitted to IEEE Computer
- [5] H.C.Lee et al.: *Grid-enabled high-throughput in silico screening against influenza A neuraminidase*, IEEE transactions on nanobioscience. 01/2007; 5(4):288-95.
- [6] J.T.Moscicki et al.: Ganga: a tool for computational-task management and easy access to Grid resources, <http://arxiv.org/abs/0902.2685>, submitted to Computer Physics Communications
- [7] N. Santos and B. Koblitz: *Metadata services on the Grid*, Nucl. Inst. & Meth. in Phys. Res. A 559(2006) pp 53-56.
- [8] N. Santos and B. Koblitz, [Metadata services on the grid](#), in Proc. of Advanced Computing and Analysis Techniques (ACAT'05), Zeuthen, Berlin, May 2005.
- [9] N. Santos and B. Koblitz, [Distributed Metadata with the AMGA Metadata Catalog](#), in the proceedings of the Workshop on Next-Generation Distributed Data Management HPDC-15, Paris, France, June 2006.
- [10] N. Santos and B. Koblitz, *Security in distributed metadata catalogue*, Concurrency and Computation: Practice & Experience 20 (2008), pp 1995-2007.
- [11] J. Montagnat, A. Frohner, D. Jouvenot, Chr. Pera, P. Kunszt, B. Koblitz, N Santos, C. Loomis, R. Texier, D. Lingrand, P. Guio, R. Brito Da Rocha, A. Sobreira de Almeida and Z. Farkas: *A Secure Grid Medical Data Manager Interfaced to the gLite Middleware*, J Grid Computing (2008) 6:45–59 DOI 10.1007/s10723-007-9088-2.