HEP SSC Work Plan (Part B)

1 HEP SSC

1.1 HEP SSC Description

General Concept and Objectives

In the past years, when grid computing emerged as the paradigm on which to base the computing systems of new High Energy Physics (HEP) experiments, specific support on grid integration provided by projects such as EDG and EGEE was essential to the successful adoption of grid computing and it is still having a major role in the development and operation of the experiment computing systems.

For this reason, the main purpose of the HEP SSC is to support the HEP experiments at CERN and elsewhere in using effectively the European Grid infrastructure.

It will allow the HEP community to exploit the scientific potential of the Large Hadron Collider (LHC) machine. All the research centres participating to the community, mostly via the Worldwide LHC Computing Grid (WLCG) project, will profit from the support provided by the HEP SSC.

HEP experiments have to collect, process and analyse typically very large quantities of data (tens of petabytes per year for the LHC alone after startup) and involve several thousands of researchers worldwide. The data processing is done on grid resources, which in the case of LHC are coordinated by the WLCG project.

The communities that will be supported by the HEP SSC are the LHC and other HEP experiments (linked by the science) as well as international ones (UNOSAT, EnviroGRID, International Telecommunications Union) – typically linked by the technology deployed.

The HEP SSC intends to exploit the well established and very successful support model developed during the last eight years by previous EU Grid projects and by WLCG and the competence of a support team which has been involved since a long time in similar activities.

More details can be found under "Impacts" below.

1.2 Call Objectives

Please describe how your SSC will address the following objectives. This should be a high-level description, so please limit the response to a couple of paragraphs for each objective. Not all SSCs need to address all of the objectives. If your SSC does not address one of them, please just write "Not applicable".

1. Deployment of e-Infrastructures in research communities in order to enable multi-disciplinary collaboration and address their specific needs.

Response:

Although the primary goal of this SSC is to support the High Energy Physics experiments at CERN and elsewhere, a number of the tools developed have already been adopted by a range of disciplines – including others in this "SSC cluster" as well as those beyond (UN initiatives and EU-funded projects such as EnviroGRIDS and PARTNER – a hadron-therapy project). Such inter-disciplinary collaboration is considered of great importance both to all partners and to the community as a whole and ways of expanding this through the Heavy User Community of EGI and beyond will be explored. This is true both "vertically" (i.e. within a given SSC) as well as "horizontally" – i.e. across distinct SSCs. (e.g. collaboration with Fusion (Ganga), Life Science (Ganga + GEANT4).

2. Deployment of end-to-end e-infrastructure services and tools, including associated interfaces and software components, in support of virtual organisations in order to integrate and increase their research capacities.

Response:

This is essentially the *raison d'être* of the proposed support centre. In particular, one of its main goals is to support the High Energy Physics and related communities at this critical phase of LHC startup and exploitation. This involves approximately 10,000 researchers worldwide who need to access and analyze data 24x7 using worldwide federated grid resources. The service and user support to this community – enabling them to maximize the scientific and discovery potential of the LHC machine and the detectors that will take data at it – is a fundamental goal.

3. Building user-configured virtual research facilities/test-beds by coalition of existing resources (e.g. sensors, instruments, networks, and computers) from diverse facilities, in order to augment the capacities of research communities for real world observation and experimentation.

Response:

In the context of WLCG, this is performed via the Memorandum of Understanding (MoU) that brings together CERN, the experiments and the resource providers around the world with day-to-day supervision and decision making performed by a WLCG Management Board consisting of members from all the WLCG project actors. This involves several distinct grids – currently EGEE, Open Science Grid (OSG) and NorduGrid – and numerous funding agencies (the signatories of the MoU). Several bodies exist to monitor that pledges and commitments are met and to review requests for additional resources and the schedule for acquiring and deploying them (Computing Resource Review board and Computing Resource Scrutiny group).

4. Addressing human, social and economic factors influencing the creation of sustainable virtual research communities as well as the take up/maintenance of e-Infrastructure services by communities.

Response:

This SSC and WLCG will be key drivers behind the interoperation of the gLite, ARC (NorduGrid) and OSG middleware stacks and related services. WLCG also has partners in Latin America and Asia Pacific.

Furthermore, one of the key challenges that faces fundamental research, such as

High Energy Physics, is to allow researchers from around the world to fully participate in their experiments – which may be physically located on the other side of the world – whilst still playing a key role in the scientific and cultural life of the University or Research Institute for which they work. Realising that education is key to the long-term success of economies and societies as a whole, ways whereby this ambitious goal can be achieved are of great importance. One of the significant advantages of grid computing as compared to previous less integrated types of remote working is the realisation of worldwide virtual research communities that can consist of thousands of researchers at hundreds of institutes where researchers are not impeded by distance and can play equal roles regardless of location. This ability has enabled LHC experiment member countries to invest in local and regional computing infrastructures at national laboratories and universities, with ten first level and over two hundred second level sites, confident that this infrastructure can be used. Success of this e-infrastructure project will reinforce this confidence leading to increased investment. Socio-economic benefits will include local employment and the continued development of local and regional centres of excellence. It has also strongly contributed to the success of worldwide distributed collaboration on grid services, whereby a highly functional data processing and analysis system can be run despite the challenges of multiple management domains, time zones, local priorities and other such challenges.

5. Integrating regional e-Infrastructures and linking them to provide access to resources on a European or global scale.

Response:

The Worldwide LHC Computing Grid (WLCG) is very much a federated grid and builds on today's EGEE infrastructure, together with grid resources provided through OSG in the US, NorduGrid in the Nordic countries as well as partners in other regions of the Americas and throughout the Asia-Pacific region. This is essential given the fully global nature of High Energy Physics and will be an important component of the proposed work.

Response for 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 Report[1] 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;

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

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 of the order of the requirements of the LHC experiments or above. A detailed evaluation is under way. As a result of the later start date 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 from 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 throughout the beam period. However 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 large number 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 Information and Communication Technologies (ICT) infrastructures and requires transversal tools across all communities.

In the long run FAIR has to become an e-infrastructure. In the shorter term the construction of FAIR must be accompanied by a raised level of e-infrastructure awareness and usage. It is therefore important to closely cooperate with WLCG. PANDA and CBM have already started using the grid for detector simulations. By the end of this project all FAIR communities will have developed and elaborated their ICT road maps.

Response for ILC:

The high energy particle physics community is planning already for projects after LHC with complementary physics programs. The **I**nternational **L**inear **C**ollider community (ILC) has started to study detector concepts in the context of EU-funded projects such as EUDET-JRA1. Grid computing is considered as an important part of the computing strategy and Grid resources have already been utilized massively for detector simulations as a proof of concept. Furthermore, the CALICE collaboration, the LCTPC and EUPixelTelescope groups used the Grid to store testbeam data.

Within this SSC the Grid e-infrastructure will be enabled for ILC by deploying community specific Grid services, implementing user and application support structures, and integrating analysis tools. Since the ILC community will certainly have - and has already - a wide overlap with the LHC experiment groups, a common e-infrastructure is an efficient way to join forces and achieve sustainability.

1.3 Interactions with Other SSCs

Please list possible interactions/collaborations with the other SSCs involved in this project.

In terms of existing communities the collaboration between HEP and Life Sciences goes back several years. It has been based on the existing set of tools developed by HEP, and used by many communities, such as metadata catalogues or simulation toolkits. Regarding the simulation toolkits, the GEANT4 package is used by life sciences in the lowest energy range provided by the toolkit. In parallel the validation of the GEANT4 toolkit has been performed within NA4 HEP and can be expected to continue within the HEP SSC, with a clear benefit to many other SSCs.

Close collaboration with Life Sciences on the existing common toolset is expected to continue. For example, through the PARTNER project – for which 3 Marie Curie doctoral students are hosted at CERN in the Grid Support group – further collaboration with Life Sciences will be required.

Please also list possible interactions/collaborations with SSCs that are NOT involved in this project.

Disciplines such as astro-particle physics and fusion have close scientific connections and it would be natural to seek collaboration and possible synergies. Such work has already been started through a number of initiatives.

In the case of the fusion community the same user analysis infrastructure used by ATLAS and LHCb has been adopted – this infrastructure has been also generalized to be applicable to any new community that wishes to use the grid. Several ITER (originally "International Thermonuclear Experimental Reactor") applications have used this infrastructure – demonstrating more general technology transfer between SSCs. Follow up on this collaboration has been agreed by the two communities.

Successful collaborations have also been setup with United Nations / EU initiatives, once again using the same tools and procedures. These collaborations open the door to the establishment of a stable UN-HEP SSC Grid platform able to add new agencies with a minimum of effort.

1.4 Partners

Please provide a list of partners that will be involved in your SSC and the necessary contact points for the partner. If a partner will participate but not receive funding from the Commission (i.e. is completely "unfunded"), please indicate that in the table. The administrative contact will be someone from the institute to contact about legal and financial issues.

Acronym	Full Name	Unfunded?	Country
CERN	European Organization for Nuclear Research	No	CH
DESY	Deutsches Elektronen Synchrotron	No	DE

GRIDPP	UK Computing for Particle Physics	No	GB
GSI	GSI Helmholtzzentrum für Schwerionenforschung GmbH	No	DE
INFN	Istituto nazionale di fisica nucleare	No	IT
UiO	University of Oslo	No	NO
FZU	Institute of Physics of the Academy of Sciences of the Czech Republic	No	CZ
OSG	Open Science Grid	Yes	USA

Acro-nym	**Full Name**	**Cou ntry**	**Scientific Contact**	**Admin. Contact**
CERN	European Organization for Nuclear Research	CH	Jamie Shiers Jamie.Shiers@cern.ch	Svetlomir Stavrev Svetlomir.Stavrev@cern.ch
DESY	Deutsches Elektronen Synchrotron	DE	Volker Guelzow; +49 40 8998 1771, volker.guelzow@desy.de	Uwe Wolframm, +49 40 8998 3870, uwe.wolframm@desy.de
GRIDPP	UK Computing for Particle Physics	GB	David Britton d.britton@physics.gla.ac.uk	David Britton d.britton@physics.gla.ac.uk
GSI	GSI Helmholtzzentrum für Schwerionenforschung GmbH	DE	Peter Malzacher, P.Malzacher@gsi.de	Johannes Heilmann, J.Heilmann@gsi.de
INFN	Istituto nazionale di fisica nucleare	IT	Claudio Grandi Claudio.Grandi@cern.ch	Giorgio Pietro Maggi Giorgio.Maggi@ba.infn.it
Oslo	University of Oslo	NO	Farid Ould-Saada, farid.ould-saada@fys.uio.no	Farid Ould-Saada, farid.ould-saada@fys.uio.no

FZU	Institute of Physics of the Academy of Sciences of the Czech Republic	CZ	Milos Lokajicek, Milos.Lokajicek@fzu.cz	Milos Lokajicek, Milos.Lokajicek@fzu.cz
OSG (un-funded)	Open Science Grid	US	Ruth Pordes, ruth@fnal.gov	Ruth Pordes, ruth@fnal.gov

1.5 Work Package HEP NA.SSC.1

1.5.1 Overview and Effort

Work Package Number	HEP NA.SSC.1
Start Date	M1
End Date	M36
Activity Type	COORD
Partner Acronym	**Effort in Person-Months**
CERN	108
Oslo	36
INFN	36
OSG (non-funded)	72

1.5.2 Objectives

- Dissemination of the progress and achievements of the e-infrastructure within the scientific and technical community and to the wider public and ensure that all deliverables produced by the project that do not include financial information or security-related issues will be made public (the project does not make claims on the IPR of the scientific results/data produced on the Grid infrastructure). This in turn will help to promote the usage of grid technology and help to establish its use in new communities and/or areas.
- Multi-directional communication with middleware providers to ensure that the needs of the community are met: European Middleware Initiative (EMI), the Advanced Resource Connector (ARC), gLite, Open Science Grid (OSG): testing and collaborative deployment of the Virtual Data Toolkit (VDT) and its components used by EGI, including Build and Test; testing and collaboration with OSG/US software collaborative developments used by EGI - VOMRS, Myproxy, MYOSG, Condor;

- Multi-directional communication with EGI operations and user support and their counterparts in other grids to ensure that the needs of the community are met (e.g. OSG through regular face to face meetings and workshops) and regions (e.g. Asia-Pacific, including the International Symposium on Grid Computing (ISGC) series of symposiums in Taiwan which brings Asia-Pacific partners face-to-face on an annual basis and in which WLCG is typically involved) (target: common and interoperable operations, architecture, policy and security work);
- Organization of regular workshops and conferences inter- and intra-VO (similar to EGEE User Fora and WLCG workshops – the latter being co-located with the Computing in High Energy Physics conferences which is the main computing-related conference for the HEP and related communities and takes place every 18 months rotating through Europe, North America and elsewhere);
- Overall WLCG Service Coordination and liaison with other WLCG structures and bodies;
- Tier2 coordination, Network coordination;
- Address the long-term sustainability of this activity: to explore how the developments made by this SSC can be supported and continue in the future.

1.5.3 Description of Work

Please provide a detail description of the work to be carried out within the work package to meet the objectives stated above. If there are multiple distinct activities, then please identify these through subtasks.

Task 1: Service coordination and liaison with other projects.

Service coordination and liaison is an on-going task that is essential to providing a world-class service and to ensure cooperation and inter-operation across widely distinct management and technical domains. It is accomplished through regular meetings, conference calls and workshops ranging from daily (for WLCG operations conference calls) to (bi-)annually for inter-operations meetings and larger (200-300 attendee) workshops. A work-plan is best described by the existing and foreseen meetings and other interactions.

Event	Recurrence	Purpose	Attendees
WLCG operations conference call	Daily		Representatives from experiments, Tier0 and Tier1 sites, major service providers (some 10-20 attendees)

WLCG workshops	3-4 times per year	Thorough analysis of top issues	100-300 attendees, depending on theme
Interoperations workshops	At least annually	Key issues regarding interoperation between different grids	10-20 people
Middleware, user support and operations	Daily	On-going issues with service deployment and delivery	Typically small focused discussions, conference calls plus strategy presentations at the above workshops

An additional activity foreseen for this task will be to organize presentations of the progress and achievements of the e-infrastructure within the wider scientific and technical community and to the broader public. This would apply to the several major international conferences per year which bring together large numbers of scientists and engineers covering a wide spectrum of activities such as the conferences on Computing in High Energy Physics (CHEP). For the wider public this involves work with the CERN press office in releasing material intended for journalists and relating to progress in this area. In the last year three grid related press releases have been made. When the LHC first started several hundred television stations worldwide participated and CERN has a high profile in the world's media so such releases have a strong impact.

The spreading of good practices, consultancy and training courses for new users are addressed through regular meetings and themed workshops – this is an on-going activity which needs to be continued, particularly during the critical early years of the LHC's operation. During these events the status of the services and the overall WLCG operations is reported and compared against the service availability expectations of the HEP communities (which in some cases can be around 99% for specific and critical services at large grid sites.) Standard operations procedures regarding service development, hardware management and maintenance have been largely discussed at several forums and are followed up on a regular basis with the grid sites that are supporting the HEP communities – constant vigilance is required to maintain the required service level. In addition, these procedures and standards have been shared with other international grid initiatives, also outside Europe. The goal of these initiatives is to spread the HEP operations requirements to other grid communities in order to establish stable collaborations as required by the corresponding user communities.

In the past the HEP community has led the creation and maintenance of grid user guides that have benefited the whole European Grid community thus contributing to dissemination of knowledge and internal / external communication. The maintenance and support of these guides as well as further introductions, FAQs and recipes will continue and will be essential as a growing number of nonexpert users turn to the grid for analysis of the data produced at the LHC. The SSC will also maintain the existing level of effort in terms of presentations, participation to Grid Forums (regional and international), tutorials and courses appropriate to the tools supported by this community such as Dashboard, Ganga, storage solutions and so forth.

Task 2: Middleware coordination

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 and will be). 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 before the annual accelerator startup. 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 and 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 during the year to provide the new versions (see above).

Task 3: Investigation and implementation of sustainability plan

For HEP and related experiments – and in particular those as long-lived as the LHC collaborations – long-term sustainability is of paramount importance. This task will investigate, recommend and start the implementation of such a plan, for which the details might vary according to area. For example, distributed analysis support is probably best handled by a small number of expert sites such as those mentioned above under the Distributed Analysis support task complemented by similar sites in other regions (the Americas, Asia-Pacific). On the other hand, integration and operations support, are key tasks that have a centre of gravity at the host laboratory of a given collaboration. Continued evolution from Integration support to Operations can *a priori* be expected.

1.5.4 Partner Contributions

Partner	Contribution
CERN	CERN is responsible for the overall WLCG service coordination: since almost two years CERN runs daily meetings to organize experiment activities together with sites, moreover it has organized regular WLCG “Collaboration” and topical workshops. Interoperation workshops are co-organized with e.g. OSG. CERN also hosts and drives central WLCG coordination bodies like the Grid Deployment Board and the Management Board. Finally, CERN has been providing, for the duration of the whole EGEE project, liaison functions between the experiments and the WLCG infrastructure.
Oslo	Grid interfacing expertise and middleware liaison. The group at the University of Oslo has several middleware and user-level grid tool experts and heads the NorduGrid and ARC activities. All components of the future EMI, of which ARC is one, will be needed and used by the supported communities, directly or through tailored interfaces and tools. The Oslo group offers its competence in this area at the level of 1 FTE.
INFN	INFN always had a leading role in the development of middleware as well as in the integration of middleware services with the application frameworks. INFN will act as a liaison to EMI (1 co-funded FTE).

1.5.5 Deliverables and Milestones

No.	WP No.	Title	Lead Partner	Nature	Due Date	Description
M.NA.HEP.1.1			CERN	O	PM06	HEP SSC workshop
M.NA.HEP.1.2			CERN	O	PM12	HEP SSC workshop
D.NA.HEP.1.1			CERN	R	PM11	Report on different options for sustainability

Due Date	D/M	Result Type	Description
PM06	M	Workshop	HEP SSC workshop
PM11	M	Workshop	HEP SSC workshop

PM11	D	Report	Report on different options for sustainability
PM??	M	Meeting	Annual interoperations meeting
PM18	M	Workshop	HEP SSC workshop
PM23	D	Report	Report on status of plan for adoption and sustainability
PM23	M	Workshop	HEP SSC workshop
PM??	M	Meeting	Annual interoperations meeting
PM30	M	Workshop	HEP SSC workshop
PM35	D	Report	Status of implementation of adoption and sustainability plan
PM36	M	Workshop	HEP SSC workshop
PM??	M	Meeting	Annual interoperations meeting

1.5.6 Risks

We must demonstrate that we've analyzed what can go wrong with our work plan and have planned contingencies if things do go wrong. Please list possible risks for the work plan (both internal and external), their effects, and mitigation strategies.

Risk	Impact	Occurrence Probability	Mitigation
Poor or lack of execution	Severe loss of service experienced	Low in most areas: this is a well understood area.	The need for good communication / coordination / liaison is well understood in most areas of WLCG over a period of several years. Events such as "collaboration workshops" are now an accepted part of our culture. Areas where this needs to be improved include network coordination, Tier2 coordination and that with sites in Asia-Pacific. These concerns are reflected accordingly in the work plan.

1.6 Work Package HEP SA.SSC.1

1.6.1 Overview and Effort

Work Package Number	HEP SA.SSC.1
Start Date	M1

End Date	M36
Activity Type	SVC
Partner Acronym	**Effort in Person-Months**
CERN	288 (= 8FTEs for 3 years)
INFN	216
GridPP	72
Oslo	36
FZU	36

1.6.2 Objectives

- User and application support services, including support for grid integration, production data processing and end-user analysis;
- Grid infrastructure / service deployment and support, including monitoring of resource usage and service availability / reliability, service coordination, debugging of complex middleware service problems and feedback to service / middleware providers;

1.6.3 Description of Work

Please provide a detail description of the work to be carried out within the work package to meet the objectives stated above. If there are multiple distinct activities, then please identify these through subtasks.

The core of this work package is the support for the associated communities and their respective production and analysis activities. As these activities are somewhat different in nature, they are presented as separate sub-tasks.

Task 1: Integration Support
HEP experiments have developed elaborate computing frameworks on top of the grid middleware(s) which now operate in full production. However, the experiments will need to adapt their infrastructures to exploit new middleware functionalities, cope with issues that will inevitably arise during data taking and improve the current operational model to increase automation and reduce the need for manual intervention. Therefore, effort in the area of support for the integration of grid middleware with the application layer is required. This will largely consist of:

- Testing of new middleware features and functionality in pre-production environments, as well as stress testing of key components following experiment requirements. This includes negotiation of service setups with various NGIs and middleware providers, definition of the test environment, scenarios and metrics, development of the test framework, test execution and follow up.
- Integration of experiment specific information in high level monitoring

frameworks. The 4 main LHC experiments – ALICE, ATLAS, CMS and LHCb – developed specific monitoring frameworks for both workload and data management; the aim is to provide a general view of the experiments activities oriented to different information consumers: sites, other experiments, WLCG coordination.

- Development of experiment specific plug-ins to existing frameworks. WLCG relies on complex frameworks such as Service Availability Monitoring (SAM), Service Level Status (SLS) and NAGIOS to measure site and service availability and reliability and to implement automatic notification and alarms. The experiments can benefit from a common infrastructure, developing specific plug-ins.
- Further developments oriented to integration of middleware with the application layer. This includes maintenance of end-user distributed analysis tools and frameworks and their related VO-specific plug-ins.
- Provision of a scalable and sustainable distributed support framework to support large user communities on all grid infrastructures used by a given VO.

Task 2: Operation Support

The LHC experiments are running their computing activities in production mode since many years. The WLCG infrastructure is widely distributed (more than 200 sites), heterogeneous and resources are loosely coupled. Many day by day operational tasks need grid expertise and such requests for specialized operational support will increase with the first LHC data when experiment computing models will need to react promptly to various use cases and scenarios. In such terms, the contribution of the SSC can be summarized in:

- Offer general grid expertise for identification and solution of middleware issues as well as site configuration and setup problems. This includes a possible risk analysis and definition of action plans to prevent escalation of criticality.
- Development of experiment specific operational tools. Such tools include intelligent mining of grid monitoring data (for both workload and data management), automation of workflows and procedures, enforcement of data consistency across various services (storage and catalogs).
- Support for the integration of experiment specific critical services into the WLCG infrastructure. This includes service deployment, definition of escalation procedures and support models.
- Development and operation of tools which facilitate end-to-end testing of analysis workflows, including functional testing which is integrated with SAM and stress testing to investigate site- and VO-specific bottlenecks.

Task 3: Distributed Analysis Support

Scientists have been running analysis on the WLCG distributed infrastructure since many years. In addition, experiments successfully went through several challenges to test the readiness of the infrastructure and tools to massive end user analysis. However, with the arrival of the first LHC data, chaotic access is expected to scale up by an order of magnitude and attract inexperienced Grid users. Therefore, it should not only be foreseen to have a dedicated effort for maintenance and further development of analysis tools (already accounted for in Task 1 and Task 2 above), but also a focused end-user support structure, consisting of the following activities:

- Investigation and deployment of tools which enable effective user-to-user and user-to-expert interaction.
- Coordination of support providers, namely experts from the VO user communities.
- Coordination of general and VO-specific training for end-users and support providers.

Partner Contributions

Partner	Contribution

CERN	CERN has been driving the Grid support effort for large user communities since the early days of the European Data Grid. This has been primarily focused on LHC experiments but included "gridification" of applications from other user communities (Biomed, UNOSAT and many others) – as well as other HEP communities both at CERN and at other major international laboratories around the world. The CERN Grid Support team has provided major contributions to the integration of the experiments frameworks with the WLCG infrastructure. This includes the development of interfaces for the experiment simulation and reprocessing systems, for the data management services and for the analysis frameworks. In addition, the same team has provided a considerable fraction of the manpower for the development of general and experiment specific monitoring systems (Dashboards). CERN has also developed experiment specific plug-ins to various monitoring tools such as SAM and SLS, while considerable work has been devoted to port experiment specific services into the CERN service infrastructure (installation, configuration, monitoring, procedures). From the operations point of view, CERN – in collaboration with the WLCG sites and experiments – assisted heavy user communities throughout many challenges (Service Challenges, Common Computing Readiness Challenges, Scale Test of Experiment Programme (STEP'09)) and production activities, covering all aspects of the LHC experiment computing models, with particular focus on production and analysis scenarios and data replication. Finally, many experts within the CERN Grid Support team have been and are currently covering key coordination and project management roles within the experiments. CERN will be the lead partner in all tasks described in this program of work for the WLCG community. The CERN contribution will therefore consist in the continuation and further evolution of the activities described above, with a particular focus to the support for end-user analysis.

GridPP	GridPP, the collaboration that has developed and deployed a Grid for particle physics across 17 sites in the UK, also played a vital role in launching the WLCG project as the single largest contributor of resources, funding employment of 23 FTE at CERN in the first three year phase of WLCG. GridPP continues to make major contributions to the project in the areas of operations (the GOC at RAL); security policy and operations; accounting (APEL); storage; Grid integration; production and data processing; and end-user analysis. In particular, GridPP has contributed 2 FTE to the Ganga project for most of the last 9 years, embedded in both the ATLAS and LHCb communities (and, more recently, working with other communities), initially at Oxford and Cambridge Universities but currently at Birmingham and Imperial. We request support to enable further development of Ganga as the standard Grid interface, not only for ATLAS and LHCb, but also for emerging particle physics experiments and other user communities such as Photon Science with growing requirements in grid computing. The latter is exemplified by ongoing simulation work in the HiPER and XFEL programmes. In the longer term in will be important to bridge these large scale computational activities to the routine data processing of the light source user communities. The Science and Technology Facilities Council (STFC) actively supports users across a wide range of current facilities (LCLS, FLASH, DLS, CLF) and in addition to supporting current simulation work for future light sources we will work closely with subject specific SSCs (e.g. in Materials Science and Chemistry) and broader community programs to ensure that future analysis requirements are taken into account. (Request: 3 FTE jointly funded by GridPP/STFC and the SSC)
Prague University	FZU Prague has taken part in the WLCG development from the beginning of the project contributing with computing resources to several VOs as a Tier2 centre in a state without a national Tier1. The user support is organised together with CESNET (Czech Scientific Network Provider) in the scope of EGEE activities. Prague's contribution in this area would be to develop a sustainable support model for end-user analysis at Tier3 sites. This would be done in close collaboration with the overall analysis support model developed and supported as part of the first objective. The effort required is 1 FTE.
University of Oslo	Distributed analysis on the grid with heavy experience in end-to-end work flows. Ganga expertise, as well as other grid-based analysis tools. The group at the University of Oslo has several Ganga developers and core team members with focus on interfacing Ganga to experiment or private software. Effort: 1FTE

INFN	INFN has a long history of supporting "SSC-like" activities through funding of personnel either directly attached to the experiments or else placed in the Grid Support group (or its predecessors) in CERN's IT department. The work of these people has been fundamental in adapting not only the LHC VOs computing systems to the grid but also in numerous other "gridification" projects. INFN will support the LHC experiments in particular for what concerns the integration and support of data and workload management tools with the EGI services, the development and support of service monitoring, the commissioning of data links between sites (4 co-funded FTE, 2 in each of the "Operation" and "Integration" tasks). INFN will contribute to the set up and maintenance of small distributed test beds and of test services connected to the production infrastructure to facilitate the integration of new services (or new versions of services) with the experiment frameworks (1 co-funded FTE in the "Integration" task). INFN will continue to develop and support experiments end-user analysis frameworks to facilitate the access to the EGI services, in particular CMS Remote Analysis Builder (CRAB) (1 co-funded FTE in the "Distributed analysis" task).

1.6.4 Deliverables and Milestones

No.	WP No.	Title	Lead Partner	Nature	Due Date	Description

Due Date	D/M	Result Type	Description
PM02	D	Report	Quarterly Service Report detailing Key Performance Indicators, Service and Resource Utilization, Availability and Reliability, Reports on Major Service Incidents and / or Degradations, Outlook for next period.
PM05	D	Report	Quarterly Service Report
PM08	D	Report	Quarterly Service Report
PM11	D	Report	Quarterly Service Report
PM14	D	Report	Quarterly Service Report
PM17	D	Report	Quarterly Service Report
PM20	D	Report	Quarterly Service Report

PM23	D	Report	Quarterly Service Report
PM26	D	Report	Quarterly Service Report
PM29	D	Report	Quarterly Service Report
PM32	D	Report	Quarterly Service Report
PM35	D	Report	Quarterly Service Report

1.6.5 Risks

We must demonstrate that we've analyzed what can go wrong with our work plan and have planned contingencies if things do go wrong. Please list possible risks for the work plan (both internal and external), their effects, and mitigation strategies.

Risk	Impact	Occurrence Probability	Mitigation

Staffing / contract policies	Insufficient manpower: inability to fully or efficiently exploit scientific potential of the LHC, particularly in Europe; lack of competitiveness.	High	Had the LHC not suffered a number of technical setbacks, we would now be completing the second full year of data taking, following the pilot run foreseen for late 2007. The staff reductions that are already taking place, as a result of funding lines ending (include EGEE III) and contract policies (which limit the total amount of time an individual can spend on a contract of Limited Duration), would have come at a time when the inevitable startup issues that we are still to face would hopefully have been resolved. The "Experiment Integration Support" (EIS) team, funded through a combination of EU (EGEE), INFN and CERN budget lines, has reduced from 8 FTEs at the end of 2008 to a low of 4, from which it has recently recovered slightly by the addition of one FTE. CERN intends to add further resources to this area within the limits of what is possible: hopefully 3-5 additional FTEs will be added by the first half of 2010. However, the replacement of staff with up to 6 years experience with relative newcomers is far from optimal at this critical stage. Ways of continuing at least some of the short term (maximum 3 years) staff until EU funds might become available are being investigated together with INFN, a long-term partner in this area.

Major loss or service degradation	Major disruptions to experiments' production and/or analysis. Depending on the specific computing models, analysis severely impacted in one or more regions.	High	Service problems – which can be caused by issues ranging from natural disasters such as typhoons, hurricanes and tsunamis to more mundane reasons such as construction (responsible for numerous network outages) power and cooling, hardware or software failures or misconfigurations – are simply inevitable. Indeed, "service" is measured just as much by response to problems as to the steady state of smooth running. Through a small set of light-weight operations procedures and tools we have repeatedly demonstrated our ability to cope with even the most daunting of problems. It requires, however, constant vigilance and effort – through Service Incident Reports and analyses, regular Service and Operations reports and follow-up and extensive coordination between sites, service providers and experiments. Lack of effort in this area is guaranteed to translate to numerous and all too often prolonged service problems and is hence to be avoided at all costs.

Analysis-related issues	Major disruptions or loss of efficiency in services for end-user analysis	Medium to high	This remains one of the largest unknowns in terms of service delivery to this community and for which real data taking is essential. Whereas most production activities can be scheduled in case of resource bottlenecks, this is rarely possible in the case of end-user analysis (except for specific cases, such as the use of "analysis trains", which effectively turn unscheduled, chaotic activities into scheduled, largely sequential processing). It will therefore be particularly important to have an adequately staffed analysis support team or teams that can respond to issues in this area in an agile fashion.
Failure to handle user issues at an acceptable rate (24h turn-around time)	Overload of the second-line support. This is entirely in the scientific community. The net effect is that bug correction will be slowed down as well as the production of the results	Low	Extensive prototyping has given us confidence that we have a workable solution for supporting large users communities
Unavailability/discontinuation of external tools	Need to develop replacement tools.	Low	There is an abundant choice of web based tools and we are confident that replacement components can be found. We are also selecting mainstream components which have a reasonable probability to be stable and to evolve in a non-disruptive way.

1.7. Work Package HEP SA.SSC.2

1.7.1 Overview and Effort

Work Package Number	HEP SA.SSC.2
Start Date	M1
End Date	M36
Activity Type	SVC, RTD
Partner Acronym	**Effort in Person-Months**
GSI	144

1.7.2 Objectives

- Providing a single entry point for the grid infrastructure of the FAIR experiments;
- User and application support services, especially for detector simulations and data/service challenges;
- Support for the integration of the FAIR computing framework in the grid infrastructure;
- Development of the FAIR grid computing strategy;

1.7.3 Description of Work

Please provide a detail description of the work to be carried out within the work package to meet the objectives stated above. If there are multiple distinct activities, then please identify these through subtasks.

The FairRoot framework is an object-oriented simulation, reconstruction and data analysis framework based on ROOT – developed at CERN and widely used within the HEP community and beyond, and the Virtual Monte-Carlo (VMC) interface. It includes core services for detector simulation and offline analysis. It is used by the CBM, PANDA and NUSTAR collaborations. The development of the FairRoot framework is an ongoing work. New features are developed in close collaboration with the FAIR experiments, as well as the ROOT team at CERN.

The two HEP-like FAIR experiments have already gained experience with simulations on the grid, based on AliEn, developed by the LHC experiment ALICE. The core of this work package is to support the mutually beneficial collaboration with the FairRoot team as well as grid deployment and operation. In a series of increasingly complex productions, data challenges and prototypes of multi user analysis the grid infrastructure for FAIR will be developed.

To be able to support data analysis of the distributed data sets by many users from different FAIR communities it is essential to set up a corresponding infrastructure. Users need to be made aware of the upcoming FAIR e-infrastructure as well as to be trained how to use it. In case of problems the support staff needs to be able to solve them fast and efficiently. To be able to accomplish this it is essential to make use of synergy effects in terms of the software infrastructure used for simulation, analysis, and distributed computing. Also a close collaboration between the various developer and support teams including the user communities is necessary.

Many of the main contributors to FAIR are also involved in LHC. Therefore a workforce of 1 FTE per FAIR community is sufficient.

Partner Contributions

Partner	Contribution
GSI	GSI is the host laboratory for FAIR. It has a strong record in Grid computing as a large ALICE tier-2 centre with an integrated interactive analysis facility. It is part of the distributed ROC of the Swiss/German federation of the EGEE project. (Request: 4FTE, two funded by the SSC two by GSI)

1.7.4 Deliverables and Milestones

No.	WP No.	Title	Lead Partner	Nature	Due Date	Description

Due Date	D/M	Result Type	Description
PM12	D	Report	Experience report on detector simulation of the FAIR experiments on the Grid. Definition of baseline services and technologies for FAIR Grid computing.
PM24	D	Report	Proof of concept for the FAIR Grid computing approach.
PM36	D	Report	ICT Roadmap for the FAIR communities (APPA, CBM, NUSTAR, PANDA)

1.7.5 Risks

We must demonstrate that we've analyzed what can go wrong with our work plan and have planned contingencies if things do go wrong. Please list possible risks for the work plan (both internal and external), their effects, and mitigation strategies.

Risk	Impact	Occurrence Probability	Mitigation

1.8 Work Package HEP SA.SSC.3

1.8.1 Overview and Effort

Work Package Number	HEP SA.SSC.3
Start Date	M1
End Date	M36
Activity Type	SVC
Partner Acronym	**Effort in Person-Months**
DESY	72 (co-funded)

1.8.2 Objectives

- Providing Grid core services for ILC, incl. data management;
- User and application support for ILC detector simulation and testbeam experiments related to ILC, such as CALICE, LCTPC or the EUPixelTelescope;
- Adoption of existing components and the integration of ILC community specific frameworks;

1.8.3 Description of Work

Please provide a detail description of the work to be carried out within the work package to meet the objectives stated above. If there are multiple distinct activities, then please identify these through subtasks.

Grid technology has found its way into High-Energy Particle Physics (HEP) through the vast computing demands of the LHC. It is now seen as a major compute paradigm for the LHC data analysis and it is starting to become the main computing platform for most of the other running or planned HEP experiments. The **I**nternational **L**inear **C**ollider community (ILC) will extend the e-infrastructure beyond WLCG. The ILC community has started to utilize the EGEE/WLCG Grid infrastructure for their detector studies, which are carried out in the context of the EU project EUDET as a proof of principle. The just recently validated Letter of Intend (LoI) of the ILD detector concept (a joint European/Japanese project) heavily relied on massive Grid-based detector simulations. The CALICE collaboration as well as the TPC and Pixel telescope groups of EUDET-JRA1 started to use the Grid to store and distribute testbeam data, which were recorded at DESY, FERMILAB, and CERN.

The ILC community gains access to the Grid through two global EGEE VOs which are already supported by a significant number of sites all over the world, providing massive computing resources far beyond the capabilities of a single institute. Since all collaborating sites also participate in the tier-structure of WLCG, resources are shared within the entire HEP community and make a common e-infrastructure a reality. The ILC users are starting to use the same support channels as WLCG such as GGUS with specific support units. In order to make the Grid the main platform for user analysis, the ILC community is currently considering components, which have been developed in the context of WLCG; GANGA and AMGA. ILC opts for adoption and integration of existing software rather than for new developments.

The scope of this SSC is to further extend the e-infrastructure beyond WLCG and to foster activities to make the Grid the main computing platform for ILC. This includes the running of Grid services, which are mandatory to enable the ILC on EGEE/WLCG Grid infrastructure, the support of users and their applications, and the adoption of specific software components to the needs of the ILC users. Though some of the activities have already been started, sustained services and support is a prerequisite to successfully incorporate the ILC computing in the e-infrastructure.

Subtask 1: Grid Core Services

This subtask includes the hosting of the Virtual Organizations (VO) ('calice' and 'ilc') and the operation of mandatory Grid core services. For authorization the VOMS service is deployed. Detector simulation as well as testbeam data is stored in a central repository with tape-backend, which is accessible via Grid tools (SRM). Replica management will be provided via catalogue services (LFC and AMGA). It is planned to replicate the data to selected sites in Europe and world-wide for redundancy and improved accessibility.

Subtask 2: User Support

User registration in the ILC VOs is handled via VOMRS. User support is provided via mailing lists, which are connected to GGUS. For ILC a distinct support unit is implemented into the ROC DECH GGUS portal. Within this subtask services will be continued and adopted as needed.

Subtask 3: Adoption and integration

The ILC community has started to adopt existing Grid components such as GANGA, which require further integration of the ILC data formats to be fully functional. Adopting GANGA, which is well established for LHC analyzes, allows for joint efforts between LHC and ILC. AMGA is considered as a data management tool to allow for proper handling of collections of files and data. The development of a job submission and monitoring system is planned, which allows for production-grade detector simulations on a large scale. In this subtask support and consultation for integration and developments will be provided to the ILC groups.

Partner Contributions

DESY operates an EGEE/WLCG Grid infrastructure with a full set of Grid services, which will be used to host the ILC-specific core services (VOMRS/VOMS, LFC, WMS, and BDII).

Data storage will be provided at DESY by means of a dedicated Storage Element

(SE) with a tape-backend. Data management tools will be developed in close collaboration with the CALICE members at DESY and LAL.

DESY implemented mailing lists for ILC and takes part in the ROC DECH support.

Integration of components and developments of tools will be carried out in close collaboration with the ILC software group at DESY.

Effort: 2 FTE (72PM) co-funded by DESY

1.8.4 Deliverables and Milestones

Due Date	D/M	Result Type	Description
PM12	D/M	Report	Showing that the utilization of the Grid is feasible for ILC; availability of all services; a wide range of sites provide resources.
PM24	D/M	Report	Showing production-grade integration of frameworks and components such as GANGA and AMGA.
PM36	D/M	Report	Showing that the Grid has become a regular production platform for ILC, using components such as GANGA and AMGA.

1.8.5 Risks

We must demonstrate that we've analyzed what can go wrong with our work plan and have planned contingencies if things do go wrong. Please list possible risks for the work plan (both internal and external), their effects, and mitigation strategies.

Risk	Impact	Occurrence Probability	Mitigation
The ILC Grid activities depend strongly on the existence of the (WLCG) Grid infrastructure.	The ILC community is not able to run a separate Grid infrastructure. If maintenance for WLCG is stopped, the ILC Grid activities will come to a halt.	At least in the first years of data taking it is very unlikely that the Grid technology is given up completely.	The ILC community must maintain close contact to the Grid community to pursue further developments.
Availability of support for Grid services	The ILC community is not able to maintain Grid core services.	See above.	See above.

Allocation of Grid Resources	Resource allocation has so far been done on a voluntary basis without legal commitments.	It is likely that the amount of resources will decrease until LHC data taking and usage patterns have stabilized.	Some sites have already informally committed Grid resources to ILC. Further negotiations are planned.

1.9 Work Package HEP JRA.SSC.1

1.9.1 Overview and Effort

Work Package Number	HEP JRA.SSC.1
Start Date	M1
End Date	M36
Activity Type	RTD
Partner Acronym	**Effort in Person-Months**
CERN	36
DESY	36
INFN	36

1.9.2 Objectives

- Investigation of innovative solutions for data management, targeting not only high-throughput multi-stream random-access style usage (typical of end-user analysis) but also the integration of new industry standards and solutions into end-to-end data management solutions covering catalog, file transfer and storage aspects;

1.9.3 Description of Work

Investigation of future data management technologies.

The core storage management solutions that are in use today have their roots in a different era – some 15 to 20 years ago. Since that time not only have relative costs and capacities (such as storage and network throughput) changed enormously but also the entire IT landscape. Attempts to rationalize the inevitable diversity via standards such as the Storage Resource Manager (SRM) have had debated success: if a concept does not exist in the backend it is hard to make it 'appear' via the front-end interface. Furthermore, the available implementations vary widely in their interpretation of the agreed standard, leading to additional confusion. Finally, as the individual components have been designed and implemented almost entirely independently, large opportunities for optimization and rationalization have been lost. For example, the LHC VOs deal with sets of files (depending on their computing models) which have some strong logical connection: typically the full set is treated together in various operations ranging from transfer through to data processing. However, such concepts are not implemented in the component data management solutions – even though they would allow many operations, such as bulk network transfer or retrieval from tape, to be greatly optimized. They are typically 'unpacked' – possibly by catalog lookups – handed to the subsystems one by one and then reassembled at the target system. Such operations may occur multiple times: at the source storage system, at the file transfer stage and again at the target system. Thus an investigation of the end-to-end data management problem is long overdue. This would take into account not only the advances of recent years but also take a higher level view, covering at least catalogs, data transfer and storage / access issues. Again, although of particularly pressing concern for the supported communities, the requirement is highly generic meaning that advances in this field would benefit a range of other disciplines – as has been demonstrated on numerous occasions in the past.

1.9.4 Partner Contributions

Partner	Contribution
CERN	CERN has long experience in development and operation of large scale data and storage management systems. This includes the CASTOR mass storage system (adopted at the Tier0 and 3 Tier1s) and the LCG Disk Pool Manager, the most widely adopted storage solution for WLCG Tier2s. Additionally, CERN developed essential data management services for several WLCG experiments for file transfers (FTS) and cataloging (LFC). Finally, CERN has experience in running large production clustered file systems (AFS) since many years. Therefore, CERN would also participate in the data management futures task force with 1 co-funded FTE.
DESY	DESY is the host organization for dCache.org – one of the main storage solutions in use in HEP at many of the Tier1 and Tier2 sites. It is therefore well placed to participate in the data management futures task force with 1 co-funded FTE.
INFN	INFN has developed the StoRM storage management solution, based on code originally derived from DPM. A version of StoRM supporting a mass storage back-end is currently being deployed in production for the first time. Furthermore INFN is actively involved in storage performance and functionality testing and is therefore well placed to participate in the data management futures task force with 1 co-funded FTE.

1.9.5 Deliverables and Milestones

No.	WP	Title	Lead	Nature	Due	Description

	No.		Partner		Date	

Due Date	D/M	Result Type	Description
PM11	D	Report	Report on data management issues related to analysis recommending research strategies for the immediate future.
PM23	D	Release	Prototype release of data management components addressing the concerns highlighted in the above report.
PM35	D	Release	Pre-production release of the above.

1.9.6 Risks

Risk	Impact	Occurrence Probability	Mitigation
Lack of action	Inability to exploit new technologies, inefficient use of resources, runaway operational and support costs.	Medium to low if funded.	As above. If some minimal investment is not made in this area the consequences are likely to be much higher long-term costs.

No.	WP No.	Title	Lead Partner	Nature	Due Date	Description

1.9 Additional Costs

1.10 Impacts

General Concept and Objectives

(See also the introductory HEP SSC description)

The main objectives of the HEP SSC include:

- Help research communities to integrate their computing systems with the grid by addressing their specific needs. This applies both to communities with a long experience of grid computing and to communities which are beginning to adopt it. In pursuing this objective, the HEP SSC will encourage the use of common tools and best practices to optimize the use of the infrastructure.
- Support the experiments in their data processing and analysis on the grid.

Analysis activities are going to be particularly challenging in the case of LHC when thousands of researchers will try to run their analyses on the LHC data. The role of the HEP SSC is to provide grid expertise and help to understand and solve problems. Given the extreme complexity of distributed infrastructures, this kind of operational support has a very significant impact on the ability to use them most efficiently.

- Provide support to end users. This includes providing training and documentation focused on the experiment activities, also with the objective of developing know-how inside the communities which would allow them to become more independent of external support.
- Research innovative solutions for data management. It is a widely-held opinion that data management on the grid is too complex and fragile. A significant improvement in this area would have a big impact on the operational aspects of grid computing.
- Provide service coordination and liaison (multi-directional communication to ensure that the needs of the communities are met and that information and experience is shared) among research communities, infrastructure and middleware providers. Past experience in this area has shown that this role is extremely useful to make interactions among the involved actors more effective.

A more detailed description of the specific objectives is given in the work package descriptions.

Progress beyond state-of-the-art

A strategic goal for the operation of distributed infrastructures, which are steadily growing in terms of the amount of the distributed resources and in terms of the size of the user community, is to enable where possible automation of operations procedures based on reliable proactive monitoring. While the current functionality of the monitoring infrastructure allows the detection of eventual problems and inefficiencies of the grid or of the users' computing activities, it does not always provide a clear indication of the underlying reasons and, correspondingly, in most cases human intervention is required for resolving the issues. Accumulating expertise in operating the distributed infrastructure, integration of the software and services developed by user communities with the grid middleware, automatic aggregation and analysis of monitoring statistics and incorporating this knowledge into the operations tools will result in enabling proactive monitoring and in establishing where possible automatic recovery procedures which are not currently a common part of grid operations. This will reduce the need for human interventions and will improve the efficiency of use of the distributed resources.

Other considerable improvements are foreseen in streamlining of user analysis on the distributed infrastructure. This will be achieved by further advancement of the analysis frameworks which provide an easy and reliable gateway to the Grid infrastructure for a variety of user communities and enable processing of the wide range of applications. An example of such a framework is Ganga which has been proven to be a key component to attract new user communities to the grid. The analysis frameworks will incorporate new plug-ins to implement new grid infrastructures and to adapt to business solutions such as cloud computing. Moreover, further evolution of the analysis frameworks enables the possibility to delegate user tasks to the analysis server. The server will act on the users' behalf. It will take the most optimal decisions regarding job processing, based on an integral analysis of the monitoring information. This mode of operations will considerably simplify processing of analysis tasks from the user perspective and will allow the optimization of the use of distributed resources.

Methodology

The implementation of the project objectives will follow methodologies developed over many years of experience in supporting the LHC community and other communities and which are already proven to be effective.

The proposed members of the SSC includes some that are very experienced and have a long record of successes in achieving similar objectives: by developing tools which have been widely adopted (Ganga, Dashboard, Hammercloud) by experiments, and by providing integration and operational support. These successes allowed the experiments to be able to cope with several tens of millions of jobs run by thousands of researchers actively using the grid in 2009.

The factors on which the implementation of the objectives is based include:

- A close relationship to both experiments and middleware providers and deep knowledge accumulated on the experiment computing systems, helped by the fact that many project members have a HEP background. For objectives related to tool development and grid integration this relationship is essential to understand the experiment requirements. For objectives related to operational support it will help to understand the interaction between the experiment computing system and the grid.
- Strong commitment of SSC members to work with the experiments, by participating to their internal meetings and in some cases taking part in their computing activities.
- Development of test suites and exploration of middleware functionality.
- Organization of meetings, workshops and forums among all the involved actors.