HEP SSC Work Plan (Part B)

1 HEP SSC

1.1 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 quite a wide range of disciplines – including others in this "SSC cluster" as well as those beyond (UN initiatives such as UNOSAT, EU-funded projects such as EnviroGRIDS, PARTNER etc.) 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:

N/A

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:

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 are 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 is mirrored by 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.

1.2 Interactions with Other SSCs

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

Close collaboration with Life Sciences on the existing common toolset is expected. 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.

1.3 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	Country	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		
GRIDPP		GB		
GSI	GSI Helmholtzzentrum für Schwerionenforschung GmbH	DE		
INFN	Istituto nazionale di fisica nucleare	IT		
Oslo	University of Oslo	NO		
Prague	University of Prague	CZ		
OSG (unfunded)	Open Science Grid	US		

1.4 Work Package HEP SSC.1

1.4.1 Overview and Effort

Work Package Number	HEP SSC.1
Start Date	PM1
End Date	PM36
Activity Type	SVC
Partner Acronym	**Effort in Person-Months**
CERN	288 (= 8FTEs for 3 years)

DESY	144
GSI	144
INFN	144 (through 4 CERN fellows?)
Oslo	72
Prague	36

1.4.2 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;

1.4.3 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 workpackage 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. Within each such sub-task there is also the support for communities associated with a specific accelerator centre: CERN for the LHC and various fixed-target experiments, DESY for HERA experiments and for simulations for the International Linear Collider (ILC) together with "photon science" (light source) experiments based at that laboratory as well as GSI for a new generation of heavy ion experiments "FAIR". A number of key grid resource providers and centres of excellence make up the complement of this workpackage. Although there is inevitably some substructure that reflects the above, we present the activities as just two sub-tasks. This emphasizes the collaboration, the associated exchange of information and ideas, as well as the distributed nature of the support that is required for these communities which are largely cohesive and have a long tradition of working together in such a fashion.

Task 1: (see first objective).

Task 2: (see second objective).

1.4.4 Partner Contributions

Partner	Contribution

CERN	CERN will be the lead partner in all of these aspects for the WLCG community. Its contributions therefore correspond directly to the tasks descriptions below.
DESY	Grid technology is seen as a major compute paradigm for LHC data analysis. Nevertheless Grid technology is heavily in use in other HEP experiments as well. This includes 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.), theory and in astroparticle physics experiments such as Icecube. For those experiments the major centres – e.g. DESY&INFN – host 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; 2. further develop the grid enabled storage element dCache; 3. support the scientific users in deploying grid technology. <u>Effort</u> DESY requires 2 FTE funded by the EU and will add 2 FTE from its own staff. Community: HEP Call: 1.2.3
GSI	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 about an order of magnitude higher than the requirements of the LHC experiments. 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 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 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. Effort: 4 FTEs (1 per FAIR community)
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. The foreseen contribution would be via 4 FTEs to be placed at CERN working full time on the key objectives of this workpackage.
University of	Distributed analysis on the grid, with heavy experience in end-to-

Oslo	end work flows. Ganga expertise. 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: 2 FTEs, fully funded by the University of Oslo.
Prague University	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.

1.4.5 Deliverables and Milestones

Due Date	D/M	Result Type	Description
PM02	D	Report	WLCG Quarterly Service Report
PM05	D	Report	WLCG Quarterly Service Report
PM08	D	Report	WLCG Quarterly Service Report
PM11	D	Report	WLCG Quarterly Service Report
PM14	D	Report	WLCG Quarterly Service Report
PM17	D	Report	WLCG Quarterly Service Report
PM20	D	Report	WLCG Quarterly Service Report
PM23	D	Report	WLCG Quarterly Service Report
PM26	D	Report	WLCG Quarterly Service Report
PM29	D	Report	WLCG Quarterly Service Report
PM32	D	Report	WLCG Quarterly Service Report
PM35	D	Report	WLCG Quarterly Service Report

1.4.6 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
Staffing / contract policies	Insufficient man-power: inability to fully or efficiently exploit scientific potential of the LHC, particularly	High	Had the LHC not suffered a number of technical set-backs, we would now be completing the second full year of data taking, following the pilot run foreseen

	in Europe; lack of competitiveness.		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 investigation 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, anal-	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 (re-

	ysis severely impacted in one or more regions.		sponsible 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.

1.5 Work Package HEP SSC.2

1.5.1 Overview and Effort

Work Package Number	HEP SSC.2
Start Date	PM1
End Date	PM36
Activity Type	RTD
Partner Acronym	**Effort in Person-Months**
CERN	108
DESY	36
INFN	36

1.5.2 Objectives

- Maintenance and development of middleware components required by the communities supported by this SSC and not covered by EMI, likely to include VO boxes, LFC/FTS etc.;
- 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;

1.5.3 Description of Work

Task 1: maintenance of those additional middleware components (see description below) that are required by supported communities. These components, currently part of the gLite distribution but also (FTS, LFC) deployed at non-EGEE sites (NDGF, OSG) are fundamental to the usage of grids by the LHC community in particular. The different roles that are associated to the various Tiers and the consequent massive-scale data movement are literally the lynchpin of the mapping of the experiments' computing models to the grid. Although rather stable, on-going support and maintenance is essential – not only for routine operations such as porting to new operating system releases, responding to security alerts and so forth, but also to respond to possible changes in computing models or other requirements that might arise from experience with the first extended run of the LHC and its associated production. The VO box is another such component that is essential to host the VO-specific services (covered in task SA4 of proposal against 1.2.1.2). These services are as fundamental to the experiments' production as any other service build on gLite or other middleware stacks and it is essential that they are hosted in a stable, managed environment with full monitoring, procedures and documentation (as for any other service). Again, the maintenance load is not expected to be high, just necessary. An exhaustive list of middleware components can only be provided once the final scope of the EMI proposal is clear.

This work would benefit not only the communities that are to be directly supported by this SSC but also others covered in the same proposal, such as Life Science. Indeed, it is hard to imagine any larger scale data intensive use of grid that does not require a file catalog service as a bare minimum, if not also a reliable file transfer service, such as that implemented by the FTS. Thus, in the longer term the reintegration of these well designed service-oriented components will most probably need to be considered.

Component	Description
LFC	The LCG file catalog was developed from the CASTOR storage management system and sharing its code base also with DPM – an SRM-enabled disk pool manager for sites with up to a few hundred TB of disk-only storage. It replaced the previous EDG-RLS LRC component which had both scalability and reliability problems and was preferred over the gLite FIREMAN catalog for exactly the same reasons. It is used by ATLAS and LHCb, as well as numerous other VOs within and outside the HEP community (DPM is even more widely deployed at a total of around 150 sites). For ATLAS, it is deployed as a site or "cloud" local catalog (a cloud is a Tier1 site plus its dependent Tier2s, typically in a country or region), including at non-EGEE sites such as BNL in the US (US ATLAS Tier2s also run an LFC as a local file catalog) and NDGF. LHCb uses the LFC as a global catalog but with R/O replicas at all Tier1 sites.
gLite FTS	The gLite FTS was developed in the first phase of the EGEE project and deployed in production from May 2005. It has been widely lauded for its stability and functionality at EGEE conferences and during reviews, through graphs showing multi-GB/s transfer rates sustained over many days or total transfer volumes of up to 1PB/day over several consecutive days.
gLite VO box	The gLite VO box was developed in response to a need expressed by the LHC experiments during the WLCG "baseline services" review that took place during 2005. This review established those services that were required for WLCG production. The VO box is a standard "container" for running VO-specific services, such as those identified in task TSA4.3 in the EGI proposal against 1.2.1.2. Several hundred VO boxes are deployed worldwide: ALICE requires such a box at all sites,

	whereas for ATLAS all such boxes – which host the "site services" specific to the various clouds – are run at CERN with the exception of those at BNL. VO boxes are also essential to CMS (PhEDEx et al) and LHCb (DIRAC services).
VOMS/VOMRS	[To be kept?]

Task 2: 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 relatively 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.

1.5.4 Partner Contributions

Partner	Contribution
CERN	As the original author of the cited packages, CERN would provide on-going support and maintenance for these middleware components through 2 co-funded FTEs. Based on its long experience in data and storage management, 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. It 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.

1.5.5 Deliverables and Milestones

Due Date	D/M	Result Type	Description
PM08	D	Release	Release of m/w components for that year's LHC data taking run
PM12	D	Report	Report on data management issues related to analysis recommending research strategies for the immediate future.
PM20	D	Release	Release of m/w components for that year's LHC data taking run
PM24	D	Release	Prototype release of data management components addressing the concerning highlighted in the above report.
PM32	D	Release	Release of m/w components for that year's LHC data taking run
PM36	D	Release	Pre-production release of the above.

1.5.6 Risks

Risk	Impact	Occurrence Probability	Mitigation
Lack of resources	Inability to support and maintain critical components would affect multiple VOs beyond the HEP community.	Low if funded.	The affected communities would be well advised to pool resources to provide a minimum of support should this important area be funded sub-optimally. However, this would inevitably have knock-on effects and result in (for example) poorer support for the communities.
Lack of action	Inability to exploit	Medium to	As above. If some minimal

	new technologies, inefficient use of resources, runa-way operational and support costs.	low if funded.	investment is not made in this area the consequences are likely to be much higher long-term costs.

1.6 Work Package HEP SSC.3

1.6.1 Overview and Effort

Work Package Number	HEP SSC.3
Start Date	PM1
End Date	PM36
Activity Type	COORD
Partner Acronym	**Effort in Person-Months**
CERN	108
Oslo	36
OSG (non-funded)	72

1.6.2 Objectives

- Liaison with middleware providers: EMI (ARC, gLite), OSG: testing and collaborative deployment of the 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;
- Liaison with EGI operations and user support and their counterparts in other grids (e.g. OSG) and regions (e.g. Asia-Pacific) (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);
- Overall WLCG Service Coordination;
- Tier2 coordination? Network coordination?

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.

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 above workshops

1.6.4 Partner Contributions

Partner	Contribution
CERN	CERN is responsible for the overall WLCG service coordination and has organized regular WLCG "Collaboration" and topical workshops. Interoperations workshops are co-organized with e.g. OSG.
Oslo	ARC liaison and expertise. The group at the University of Oslo heads the NorduGrid and ARC activities, and as ARC is a core part of EMI these services will be needed and used by the supported communities.

1.6.5 Deliverables and Milestones

Due Date	D/M	Result Type	Description
PM06	M	Workshop	HEP SSC workshop
PM11	M	Workshop	HEP SSC workshop
PM??	M	Meeting	Annual interoperations meeting
PM18	M	Workshop	HEP SSC workshop
PM23	M	Workshop	HEP SSC workshop
PM??	M	Meeting	Annual interoperations meeting
PM30	M	Workshop	HEP SSC workshop
PM36	M	Workshop	HEP SSC workshop
PM??	M	Meeting	Annual interoperations meeting

1.6.6 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 of 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. Things like "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 workplan.