

GridPP

Memorandum of Understanding

Among

The GridPP Collaboration as represented by the Project Leader and Project Management Board (hereinafter referred to as “GridPP”)

and

The UK Tier-1 and Grid Operations Centre, based at the Rutherford Appleton Laboratory (RAL)

and

The London Regional Tier-2, currently comprising Brunel University, Imperial College London, Queen Mary, University of London, Royal Holloway, University of London and University College London (hereinafter referred to as “London”)

and

The Northern Regional Tier-2, currently comprising Lancaster University, University of Liverpool, University of Manchester and University of Sheffield (hereinafter referred to as “NorthGrid”)

and

The Southern Regional Tier-2, currently comprising University of Birmingham, University of Bristol, University of Cambridge, Oxford University and RAL-PPD (hereinafter referred to as “SouthGrid”)

and

The ScotGrid Regional Tier-2, currently comprising Durham University, University of Edinburgh and University of Glasgow (hereinafter referred to as “ScotGrid”).

(hereinafter referred to collectively as the "Tiers" and together with GridPP as the “Parties” and each of them being a “Party”)

WHEREAS

- A. The GridPP Collaboration submitted a proposal ("The GridPP3 Proposal") to the Physics and Astronomy Research Council (PPARC) for funding to build a UK Production Grid for Particle Physics, involving the construction of a Tier-1 UK Regional Centre, support of four UK Regional Tier-2s, support for the Worldwide Large Hadron Collider Computing Grid Project (WLCG), development of Middleware, Networking and Security, support for the development of Particle Physics Applications and participation in the EU Enabling Grids for e-ScienceE (EGEE) Project;
- B. The Science and Technology Facilities Council (STFC) has approved the project and has undertaken to provide funding for the project from 1st April 2008 till 31st March 2011;
- C. The UK Institutes involved in the GridPP Collaboration have substantial computing resources that they will contribute towards the UK Production Grid in return for additional hardware and manpower support with which to operate those resources funded by STFC;
- D. The parties wish to document the structure and administrative arrangements thereof.

NOW THEREFORE THE PARTIES HEREBY AGREE AS FOLLOWS:

1. Definitions

1.1. In this Memorandum of Understanding ("MoU"):

- "The Grid" means the UK Grid for Particle Physics as defined by the GridPP Project.
- "CB" means the GridPP Collaboration Board.
- "PMB" means the GridPP Project Management Board.
- "Deployment Board" means the GridPP Deployment Board.
- "LHC" means the Large Hadron Collider at CERN.
- "WLCG" means the Worldwide LHC Computing Grid.
- "EGEE" means the EU Enabling Grids for e-ScienceE Project.
- "KSI2K" means a measure of CPU power in 1000s of SpecInt 2000 CPU equivalents.
- "TB" means useable disk storage space in TeraBytes.
- "VO" means Virtual Organisation, an abstract entity grouping Users, Institutions and Resources (if any) in the same administrative domain, e.g. a particle physics collaboration.
- "Worker Node" means the computers at the institutes where the jobs actually run excluding gateway and server machines.
- "Front Ends" means the standard set of computers provided by GridPP to each institute to act as gateways between the institutes and the rest of the Grid.
- "Target Shares" means the minimum share of a particular resource to be made available via the Grid to an approved VO in a specified time frame.

2. Purpose and Duration

- 2.1. The purpose of this MoU is to record the arrangements for the support, availability and use of the resources at the Tiers by GridPP.
- 2.2. This MoU is not binding on the Parties, it being understood however that the Parties recognise that the success of the GridPP Collaboration depends on all Members adhering to its provisions.
- 2.3. The Parties shall have no liability towards each other for any loss or damage resulting from their use of GridPP computing resources, including their participation in EGEE and WLCG.
- 2.4. The Parties acknowledge that the terms of this MoU may be varied by unanimous written agreement only.
- 2.5. This MoU shall enter into force at the date of the last signature of all the parties and shall, subject to the provisions of Clause 14, continue in force until 31st March 2011.
- 2.6. The Annexes are an integral part of this MoU.
- 2.7. The parties shall review, and as necessary update, the Annexes annually to reflect the situation on the 1st April of the current year.
- 2.8. The Parties agree that disputes arising from this MoU shall be handled, in the first instance by the PMB, with ultimate recourse to the CB.

3. The Deployment Board

- 3.1. The Parties agree that a Deployment Board, as described in the GridPP3 Proposal, shall be established to act as a forum for discussion amongst the Parties.
- 3.2. The Terms of Reference (ToR) of the Deployment Board shall be agreed by the Deployment Board and approved by the PMB and posted on the GridPP Website. The ToR shall be updated as necessary with the agreement of the Deployment Board and the approval of the PMB.
- 3.3. Each Tier shall nominate one person to represent the Tier on the Deployment Board.
- 3.4. In addition the Tier-1 Manager and the Tier-2 Technical Coordinators shall be members of the Deployment Board.
- 3.5. The GridPP Production Manager shall also be a member of the Deployment Board.
- 3.6. The Chair of the GridPP User Board and three additional user representatives from ATLAS, CMS and LHCb shall also be members.
- 3.7. Other persons shall be co-opted onto the Deployment Board as required with the agreement of all the Parties.
- 3.8. The Deployment Board shall be chaired by the Deployment Board Chair, appointed by the GridPP Project Leader.
- 3.9. A meeting of the Deployment Board shall be declared quorate for voting purposes when at least one representative of each Tier, or their deputies, plus the Chair or his deputy, are present at any scheduled meeting (in person, by phone or by videoconference) for which proper notice shall have been given.
- 3.10. Decisions of the Deployment Board shall normally be by unanimous agreement between the five Tiers and GridPP, as represented by the Chair or his deputy.

- 3.11. Decisions that cannot be agreed by the Deployment Board shall be referred, in the first instance, to the PMB for resolution.
- 3.12. The Deployment Board shall report to the GridPP PMB through its Chair.

4. Support Staff

- 4.1. GridPP agrees to support a number of Support Staff to be deployed at the Operations Centre as well as the Tiers to support the hardware and software at the institutes that is being contributed towards the Grid. The allocations are given in Annexe A.
- 4.2. The allocation of Support Staff has been agreed in the GridPP3 Proposal and ratified by the STFC through its review structures.
- 4.3. The Parties agree to these allocations.
- 4.4. If further posts become available, the Parties agree that this shall be reviewed based on past performance.
- 4.5. The allocations of the Support Staff to the institutes have been agreed within the Tier-2s themselves and any changes shall be notified to the Deployment Board Chair.
- 4.6. The Support Staff shall produce Quarterly Reports on their activities which shall be passed to the Deployment Board via the Production Manager and reported to the GridPP Project Manager through the Deployment Board Chair.
- 4.7. The Quarterly Reports shall conform to the proforma and guidance issued by the Project Manager.

5. Hardware Resources

- 5.1. GridPP agrees to provide resources for hardware to be deployed at the Tiers that is being contributed towards the Grid. The allocations are given in Annexe B.
- 5.2. The Parties agree that CPU and Disk Space allocated to their Tiers will be made available via the Grid.
- 5.3. The Tiers agree to make their CPU and Disk Data Storage resources plus their associated infrastructure and software, available on the Grid, subject to Security policies as described in Clause 11. The definition of available is described in Clause 6.
- 5.4. The parties agree that future allocations will be based upon past performance.
- 5.5. The Tier-2s shall report resource availability and usage to the Deployment Board and complete WLCG resource reports for the forthcoming year, as required.

6. Availability of Resources

- 6.1. The Tiers agree to make their resources available subject to a service level agreement agreed with GridPP. These levels of service are given in Annexe C and differ between the Tier-1 and Tier-2s.
- 6.2. Future updates to the Service Level Agreement shall be agreed by the Deployment Board.
- 6.3. The minimum response time refers to the time taken for a fault ticket against the site to be acknowledged either by the site system administrators themselves, authorised system administrators at another site within the same Tier, the Tier Technical Coordinator or any other authorised person. There is no explicit requirement on the time taken to rectify a fault other than its impact on the site's overall availability for that period.

- 6.4. The Tiers agree to provide a suitable environment for their supported VOs as agreed with the Deployment Board.
- 6.5. The Parties agree that it is the responsibility of the experiments to install and maintain experiment-specific software. The exception is that of VO boxes at the Tiers, in particular the Tier-1, where an SLA defines a detailed specification for experiment-specific services that have shared responsibility and have been globally agreed.

7. Monitoring and Accounting of Hardware Resources

- 7.1. GridPP agrees to provide monitoring and accounting software to the Tiers subject to the agreement of the Deployment Board.
- 7.2. The Tiers agree to run this monitoring, subject to Security policies as described in Clause 11 such that the performance and resource usage can be collated by GridPP.
- 7.3. The Parties agree that monitoring information from each Worker Node is not required and only aggregated results from each institute need be made available.
- 7.4. The Parties agree that the results from the monitoring and accounting software be made publicly available via the Web.

8. Target Shares and Storage Allocations

- 8.1. The Parties agree that the overall Target Shares and storage allocations for the Tier-1 Centre and total for the Tier-2s shall be set by the GridPP User Board.
- 8.2. The Parties agree that the fraction of the Target Shares and storage allocations at each of the Tier-2s shall be agreed by the Deployment Board.
- 8.3. The Parties agree that the fraction of the Target Shares and storage allocations at each Institute shall be decided within each Tier-2.

9. Software

- 9.1. GridPP agrees to provide a set of Grid software releases as defined by the Deployment Board.
- 9.2. The Tiers agree to implement and update this software as advised by the GridPP Production Manager.
- 9.3. The Parties agree that the timescale for implementation of any major release should be set by the Deployment Board and that the Tiers agree to implement the software on that timescale.

10. Network Connectivity

- 10.1. GridPP agrees to provide Network Monitoring systems so that the connectivity between each institute and the Tier-1 Centre can be monitored.
- 10.2. The Tiers agree to run this monitoring, subject to Security policies as described in Clause 11.

11. Security Policy

- 11.1. The Parties agree to implement a Security policy.
- 11.2. The policy will be defined in a set of documents. These will be available on the Deployment Board area of the GridPP Website.

- 11.3. The Security policy is intended to be identical to the policy developed for and approved by EGEE.
- 11.4. GridPP input to the approval process of updates to the EGEE Security policy shall be agreed by the Deployment Board.
- 11.5. In cases where the Deployment Board does not agree to implement an approved EGEE security policy, this fact will be recorded on the GridPP Website and communicated to the EGEE security officer.
- 11.6. The Parties agree that access to resources at the institutes can be temporarily suspended at any time without notice in case of security breaches or non-compliance with the Security policy.

12. Operations and User Support

- 12.1. The Tiers agree to participate in UK-wide Grid Operations with other relevant parties, sharing duties that are agreed through the PMB for operation of the UK Grid, WLCG, and interoperation with the rest of EGEE.
- 12.2. The Tiers agree to participate in VO-wide User Support with other relevant parties, within the VOs, sharing duties that are agreed through the VOs and the UB for user support on the UK Grid.
- 12.3. Participants in the Operations Team will comprise, at least, the Production Manager, the Operations Centre Manager, the Tier-2 Coordinators, the Security Officer, the Documentation Officer and other relevant experts who may be co-opted by the PMB.
- 12.4. Operations and User Support duties include site monitoring and reporting, EGEE Grid Operator on Duty, helpdesk triage and metrics evaluation. The associated service levels are given in Annexe C.
- 12.5. Sites agree to cooperate with the Operations Team, responding promptly to trouble tickets and other queries.

13. Management

- 13.1. The Tier-1 agrees to hold regular management meetings.
- 13.2. The Tier-2s agree to coordinate the institutes within their Tier and to hold regular management meetings.
- 13.3. The Tiers agree to report on their activities to the Deployment Board.
- 13.4. The Parties agree that the Tier activities shall also be reported to the PMB and to the STFC Oversight Committee as appropriate.
- 13.5. The Tiers agree to maintain up-to-date websites with general information about the Tier, and monitoring and other information as agreed by the Deployment Board.
- 13.6. The Tiers agree to report site problems to GridPP via the Production Manager, Deployment Board or PMB as appropriate.
- 13.7. The Parties agree that if the individual Tier Management Boards are unable to resolve a GridPP-related problem, the problem shall be referred in the first instance to the PMB with ultimate recourse to the CB.

14. Extension

- 14.1. The initial period of validity of this MoU governs the period of the GridPP3 Proposal.

14.2. The validity of this MoU shall be extended, each time for a successive period of three years beyond the initial period or, as the case may be, the previous three-year period, except as the GridPP CB may determine otherwise. This provision notwithstanding, this MoU shall terminate if and when the GridPP programme is declared closed by STFC.

15. Termination

15.1. Any Party shall be entitled to terminate its participation in this MoU on giving to the Parties at least twelve months' prior written notice of its intention so to terminate.

15.2. In the event of such termination, the STFC-funded Hardware and the Support posts at that site may be reallocated as agreed by the PMB and ratified by the CB.

IN WITNESS WHEREOF these presents consisting of this and the preceding pages are signed by the Parties as follows:

Signed for and on behalf of GRIDPP by:

Position:

Date:

Signed for and on behalf of UK TIER-1 and GRID OPERATIONS CENTRE by:

Position:

Date:

Signed for and on behalf of LONDON by:

Position:

Date:

Signed for and on behalf of NORTHGRID by:

Position:

Date:

Signed for and on behalf of SCOTGRID by:

Position:

Date:

Signed for and on behalf of SOUTHGRID by:

Position:

Date:

Annexe A: Staffing

Tier-1 Hardware Support Staff

A brief description of the service breakdown meeting the service level requirements is as follows:

CPU Service	At the hardware level the CPU resources will be run as a single large cluster with a single infrastructure for operations, installation, monitoring, and resource scheduling. At the resource scheduling level the cluster can be subdivided for periods of time to provide the flexibility to meet differing requirements of the experiments. Staff will be responsible for the security, integrity and performance of the operating system.
Data Services	10,000-20,000 disk drives will be required to provide the required storage capacity. High quality data services are the foremost requirement of a Tier-1 Centre. As well as large capacity, the service must be reliable, robust, well-managed and guarantee the integrity of the data. The tape robot is controlled by the CASTOR2 control system, which provides HSM capability needed to manage the temporary staging of data between tape and disk.
Core Services	Consoles, user file systems, monitoring, software development/deployment and conditions databases.
Operations	Machine room environment (safety, power, cooling, deliveries, installation), hardware diagnostics and repair, automation and monitoring, deploying Grid fabric management tools, tape movement and robot intervention.
Incident Response Unit	An Incident Response Unit will be developed to meet the MoU service requirements. Flexible time working will enable system staff cover from 0800-1800 Mon-Fri excluding UK public holidays. With suitable automation it is expected that the services will be kept running 24/7x365. The automation procedure will call out staff outside these hours including weekends and public holidays to diagnose problems that cannot be fixed automatically. Allowances are payable to staff who are "on-call".
Networking	Administering the centre's LAN and a share of local WAN access.
Deployment	The Tier-1 Centre will play a leading role in UK deployment, managed as part of the Grid Deployment Team to provide support for the whole of GridPP.
Experiment Support	This is outward-looking Tier-1 effort that is intended to complement dedicated experiment software support.
Management	Includes both internal management of the Centre and its team and finances, and the outward-facing interactions with WLCG, EGEE, and the other Tier centres. The existing management team has a proven record of highly efficient National and International service delivery geared to the requirements of particle physics.

Manpower will be allocated to the following sets of identified services as follows:

Work Area [FTEs]	Tier-1		
	Funding Source:	GridPP	e-Science
CPU		2.0	0.0
Disk		3.0	0.0
CASTOR Core		2.0	1.3
Core Services		1.0	0.5
Operations		3.0	0.8
Incident Response Unit		2.0	0.0
Networking		0.0	0.5
Deployment		1.5	0.0
Experiments		1.5	0.0
Tier-1 Management		1.0	0.3
Totals		17.0	3.4

Longer-term changes to the above planning will be agreed by the Deployment Board in order to best meet the required overall service levels.

The staff cover provided for the overall services described above is as follows.

- Using flexible time working, there will be system staff cover from 0800-1800 Mon-Fri excluding UK public holidays.
- Wherever possible, out of hours fault recovery will be automatic or handled by resilience within the infrastructure. On-call staff will be called, where automatic fault correction is impossible or unsuccessful. The interval allowed between fault notification and staff response will be set at a level consistent with the service level commitments. No cover will be provided over Christmas Day and Boxing Day except to respond and make safe equipment in the event of machine room emergencies such as loss of cooling. Allowances are payable to staff who are “on-call”. Further payments are made in the event of a response to a callout.

Tier-2 Hardware Support Staff

The Tier-2 allocations, based on past performance and current capacity, are

London	4.00 FTE
NorthGrid	4.50 FTE
ScotGrid	1.75 FTE
SouthGrid	2.75 FTE

with individual site allocations

Brunel	0.50 FTE
Imperial	1.25 FTE
QMUL	1.25 FTE
RHUL	0.50 FTE
UCL	0.50 FTE
Lancaster	1.25 FTE
Liverpool	0.75 FTE
Manchester	2.00 FTE
Sheffield	0.50 FTE
Durham	0.25 FTE
Edinburgh	0.50 FTE
Glasgow	1.00 FTE
Birmingham	1.00 FTE
Bristol	0.45 FTE
Cambridge	0.45 FTE
Oxford	0.20 FTE
RAL PPD	0.65 FTE

Operations Support Staff

1.5 FTEs are allocated towards monitoring, accounting and co-ordination within the UK Grid Operations Centre.

Annexe B: Hardware

Tier-1 Hardware

The Parties agree that the minimum amount of CPU and disk space made available (as defined in clause 6) shall be:

	CPU (KSI2K)				Disk (TB)			
	April 2008	April 2009	April 2010	April 2011	April 2008	April 2009	April 2010	April 2011
Tier-1	3716	6313	9313	13367	2011	3484	6044	8345

	Tape (TB)			
	April 2008	April 2009	April 2010	April 2011
Tier-1	2125	4368	7181	10432

Tier-2 Hardware

The Parties agree that the minimum amount of CPU and disk space made available (as defined in clause 6) shall be:

	CPU (KSI2K)				Disk (TB)			
	April 2008	April 2009	April 2010	April 2011	April 2008	April 2009	April 2010	April 2011
LondonGrid	2049	2859	3534	4424	340	622	1085	1540
NorthGrid	1556	2099	2675	3515	415	691	1128	1582
ScotGrid	1107	2105	2932	3532	289	487	811	1136
SouthGrid	1399	2172	2945	3624	301	515	866	1205
Total	6110	9235	12086	15096	1345	2315	3889	5463

Note: the shaded cells are nominal allocations extrapolating the formula from the first allocation. The actual allocations will be determined in May 2009 using past performance up to that point.

The final Tier-2 hardware is allocated using an agreed formula, in two tranches, one in October 2007 using metrics from the previous 6 months (07Q2, 07Q3) and one in May 2009 using metrics from the previous year (08Q2, 08Q3, 08Q4, 09Q1). The metric is currently based on SAM Test Efficiency x (CPU Delivered + Disk Available).

Units of CPU and storage (currently kSI2k and TB) may be modified according to standards agreed at the WLCG Management Board (MB) and applied uniformly across all UK sites after agreement at the GridPP PMB.

Annexe C: Service Level Agreement

The following is based upon the WLCG MoU as signed by PPARC on 17th March 2006, following discussions between PPARC, GridPP and WLCG management and was incorporated in the GridPP3 proposal submitted on 13th July 2006. It describes the services to be provided at the UK Tier-1 and Tier-2 institutes. Service levels are defined at a very coarse level. Detailed service definitions with key metrics are elaborated and maintained by the operational boards of the WLCG and GridPP projects.

Grid Services

The Tiers must provide sufficient computational and storage resources to support the proper operation of its services. In addition, the site must have sufficient network bandwidth to successfully, and continuously, pass service availability and reliability tests.

VO Support

The site must support the “ops” VO required for service availability and reliability tests. Each Tier must continuously support those user-community Experiments for which hardware is allocated. In addition, the Tiers are encouraged to support as many VOs as they reasonably can.

Tier-1 Services

The UK Tier-1 forms an integral part of the central data handling service of the Experiments. It is thus essential that each such centre undertakes to provide its services on a long-term basis (initially at least 5 years) and to make its best efforts to upgrade its installations steadily in order to keep pace with the expected growth of LHC data volumes and analysis activities.

The Experiments depend on the Tier-1 operating with a high level of reliability and availability, and consequently these services shall be provided with a high level of availability and rapid responsiveness to problems.

The following services shall be provided by the UK Tier-1 in respect of the Experiments that they serve, according to policies decided by these Experiments:

- acceptance of raw, processed and simulated data from the WLCG Tier-0 Centre, keeping up with data acquisition;
- recording and maintenance of raw and processed data on permanent mass storage;
- provision of managed disk storage providing permanent and temporary data storage for files and databases;
- operation of a data-intensive analysis facility for the LHCb experiment;
- provision of other services according to agreed Experiment requirements;
- provision of high-capacity network services for data exchange with the Tier-0 Centre, as part of an overall plan agreed amongst the Experiments, Tier-1 and Tier-0 Centres;
- provision of network services for data exchange with Tier-1 and selected Tier-2 Centres, as part of an overall plan agreed amongst the Experiments, Tier-1 and Tier-2 Centres;
- administration of databases required by Experiments at Tier-1 Centres.

All storage and computational services shall be “grid enabled” according to standards agreed between the LHC Experiments and the regional centres.

The following parameters define the minimum levels of service:

Service	Maximum delay in responding to operational problems			Average availability measured on an annual basis	
	Service interruption	Degradation of the capacity of the service by more than 50%	Degradation of the capacity of the service by more than 20%	During accelerator operation	At all other times
Acceptance of data from the Tier-0	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 during accelerator operation	12 hours	24 hours	48 hours	98%	n/a
Data-intensive analysis services, including networking to Tier-0, Tier-1 centres	24 hours	48 hours	48 hours	98%	98%
All other services – prime service hours ¹	2 hour	2 hour	4 hours	98%	98%
All other services – other times	24 hours	48 hours	48 hours	97%	97%

The response times in the above table refer only to the maximum delay before action is taken to repair the problem. The mean time to repair is also a very important factor that is only covered in this table indirectly through the availability targets. All of these parameters will require an adequate level of staffing of the services, including on-call coverage outside of prime shift.

Tier-2 Services

The following services shall be provided by each of the institutes within the UK Tier-2s in respect of the Experiments that they serve, according to policies decided by these Experiments:

- provision of managed disk storage providing permanent and/or temporary data storage for files and databases;
- operation of an end-user analysis facility;
- provision of other services, such as simulation, according to agreed Experiment requirements;
- provision of network services for data exchange with the Tier-1 Centres, as part of an overall plan agreed between the Experiments and the Tier-1 Centres concerned.

All storage and computational services shall be “grid enabled” according to standards agreed between GridPP and the Experiments.

The following parameters define the minimum levels of service per institute:

Service	Maximum delay in responding to operational problems		Average availability measured on an annual basis
	Prime time	Other periods	
End-user analysis facility	2 hours	72 hours	95%
Other services	12 hours	72 hours	95%

¹ Prime service hours are 08:00-18:00 during the working week of the centre, except public holidays.

Operations Services

This section refers to services required for the operation and management of the UK Grid.

- UK Grid Operations is a distributed activity with duties being shared between GridPP, the National Grid Service, and EGEE.
- It is responsible for maintaining configuration databases, operating the monitoring infrastructure, pro-active fault and performance monitoring, provision of accounting information, and other services that may be agreed.
- UK Grid Operations shall provide on-line coverage during prime-time hours. The service levels are defined as:

<i>Scope of the service</i>	<i>Monitoring and Accounting for UK Grid</i>
<i>Availability of monitoring services during prime-time period</i>	99%
<i>Availability during other periods</i>	97%

User Support

First level (end-user) helpdesks and Grid call centres are operated by the Experiments and the operations centre. Given that end-user support is of fundamental importance in ensuring that the Grid works the Parties agree to provide support and assistance, as required.