



August 25 2005, Jamie Shiers

Planning for SC4 and the Initial LHC Service: Step 1

Executive Summary

This document lists the initial steps that need to be completed to prepare for Service Challenge 4 and the Initial LHC Computing Service.

This document will be complemented by more detailed planning papers, covering the list of Grid Middleware components and versions required, the corresponding services and deployment strategies, operations and user support plans as well as applications area-related services.

Furthermore, experiment-specific software and components are not covered by this document.

Experiment-specific plans for SC4 will be requested and reviewed through the Grid Deployment Board / Project Execution Board as appropriate and made available via the LCG [Service Challenge Wiki](#).

To avoid changes to the numbering below, any corrections, additions or other updates will be made at the end of the document and indicated as shown below.

0. Sample correction.

Phases of Service Challenge 4

1. SC4 consists of a Setup Phase starting on 1st April 2006, during which a number of Throughput tests will be performed (see below for specific goals), followed by a Service Phase from 1st May 2006 until the 30th September 2006.
2. All service components for SC4 need to be delivered ready for production by the 31st January 2006. The list of these components and their version is given below will be detailed in a subsequent note.
3. Final testing and integration of these components and services must be completed by 31st March 2006.
4. The impact of the above schedule on compiler and operating system versions must be discussed and agreed at the Project Execution Board (PEB).

The Initial LHC Computing Service

5. The initial LHC Computing Service must be in operation by September 30th 2006, capable of handling the full target data rate between CERN and the Tier1 sites.
6. The LHC Computing Service is commissioned on 1st April 2007, via the successful completion of a series of performance, throughput and reliability tests that demonstrate readiness to operate continuously at the full target data rate, as well as at twice this data rate for sustained periods.

Timeline

7. January 31st 2006: all required components deployed at participating sites fully production ready.



August 25 2005, Jamie Shiers

8. February 1st 2006: start of integration testing of components and sites.
9. April 1st 2006: setup phase, during which site readiness is validated by a throughput test, demonstrating the nominal rate to each site. The data must be written to tape through an agreed SRM interface.
10. September 30th 2006: Start of Initial LHC Computing Service.
11. April 1st 2007: LHC Computing Service fully commissioned.

List of Sites

12. All official Tier1 sites are required to take part in Service Challenge 4.
13. All Tier1 sites must support at least 5 named Tier2 sites throughout all phases of Service Challenge 4.
14. At least 40 Tier2 sites must also participate. This includes negotiating a formal support agreement with an appropriate Tier1 site for archiving of Monte Carlo products, delivery of analysis data, the provision of the File Transfer service and the archival of backup copies of Tier2 databases if required.
15. The list of Tier1 and Tier2 sites, as well as the experiments that they support will be maintained on the Service Challenge Wiki, validated via the agreed Site Functional Tests.

Site Responsibilities

16. Site responsibilities for Services and Resources are those defined in the Worldwide LHC Computing Grid Collaboration [Memorandum of Understanding](#) (annexes 3.1, 3.2 and 3.3, reproduced in draft form below).
17. Transfer rates that must be supported from the Tier0 to each Tier1 are defined in the [LCG TDR](#), based on the nominal data rates expected during normal proton-proton running of the LHC.
18. Sites shall prepare detailed plans regarding the above two points, which shall be regularly reviewed by a committee appointed by the LCG Project Leader chaired by the LCG Service (Challenge) Coordinator and reporting to the PEB.
19. Nominal transfer rates between other sites (Tier1 to Tier2, Tier2 to Tier1, Tier1 to Tier1 and Tier1 to Tier0) will be calculated based on the Computing Models and Use Cases defined in the experiments' Computing TDRs.

Corrections and Updates

Date	Heading	Details
00/00/00	Executive Summary	Dummy text.



August 25 2005, Jamie Shiers

Annex 3. Minimal Service Levels to qualify for membership of the LHC Computing Grid Collaboration

This Annex describes the qualitative aspects of the levels of service to be provided by the Host Laboratory (CERN), Tier1 Centres and Tier2 Centres in order to fulfil their obligations as Parties to this MoU. Also described are the qualitative aspects of Grid Operations Services that some of the Parties will provide. The quantitative aspects of all of these services are described for each Party in **Error! Reference source not found..** Only the fundamental aspects of service levels are defined here. Detailed service definitions with key metrics will be elaborated and maintained by the operational boards of the Collaboration.

Annex 3.1. Host Laboratory Services

The Host Laboratory shall supply the following services in support of the offline computing systems of all of the LHC Experiments according to their computing models.

- Operation of the Tier0 facility providing:
 - high bandwidth network connectivity from the experimental area to the offline computing facility (the networking within the experimental area shall be the responsibility of each Experiment);
 - recording and permanent storage in a mass storage system of one copy of the raw data maintained throughout the lifetime of the Experiment;
 - distribution of an agreed share of the raw data to each Tier1 Centre, in-line with data acquisition;
 - one (the first) pass of full event calibration and reconstruction (in the case of pp data, in-line with the data acquisition), storage of the resulting data on tape and disk, distribution of the reconstructed data to Tier1 centres;
 - services for the storage and distribution of current versions of data that are central to the offline operation of the Experiments, according to policies to be agreed with the Experiments.
- Operation of a high performance, data-intensive analysis facility with the functionality of a combined Tier1 and Tier2 Centre, except that it does not offer permanent storage of back-up copies of data. In particular, its services include:
 - data-intensive analysis, including high performance access to the current versions of the Experiments' real and simulated datasets;
 - end-user analysis.
- Support the termination of high speed network connections by all Tier1 and major Tier2 Centres as requested.
- Coordination of the overall network design between the Host Laboratory, Tier1 and Tier2 Centres, in collaboration with national research networks and international research networking organisations.
- Tools, libraries and infrastructure in support of application program development and maintenance.



August 25 2005, Jamie Shiers

- Basic services for the support of standard¹ physics “desktop” systems used by members of the LHC Collaborations resident at CERN (e.g. mail services, home directory servers, web services, help desk).
- Administration of databases used for the storage of physics data and associated meta-data.
- Infrastructure for the administration of the Virtual Organisation (VO) associated with each Experiment.
- Provision of the following services for Grid Coordination and Operation:
 - Overall management and coordination of the LHC grid - ensuring an effective management structure for grid coordination and operation (e.g. policy and strategy coordination, security, resource planning, daily operation,...);
 - integration, certification and distribution of software required for grid operation;
 - organisation of adequate support for this software, generally by negotiating agreements with other organisations;
 - participation in the grid operations management by providing an engineer in charge of daily operation one week in four (this service is shared with three or more other institutes providing amongst them 52-week coverage).

The following parameters define the minimum levels of service:

<i>Service</i>	<i>Maximum delay in responding to operational problems</i>			<i>Average availability² measured on an annual basis</i>	
	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
Raw data recording	4 hours	6 hours	6 hours	99%	n/a
Event reconstruction or distribution of data to Tier-1 Centres	6 hours	6 hours	12 hours	99%	n/a
Networking service to Tier-1 Centres during accelerator operation	6 hours	6 hours	12 hours	99%	n/a
All other Tier-0 services	12 hours	24 hours	48 hours	98%	98%
Networking service to Tier-1 Centres	12 hours	24 hours	48 hours	98%	98%
Operations management (engineer in charge)	4 hours	6 hours	12 hours	n/a	n/a
All other services – prime service hours ³	1 hour	1 hour	4 hours	98%	98%
All other services – other times	12 hours	24 hours	48 hours	97%	97%

¹ The standard supported desktop systems are agreed from time to time between CERN and its user community.

² (time running)/(time scheduled to run)

³ Prime service hours for the Host Laboratory: 08:00-18:00 in the time zone of the Host Laboratory, Monday-Friday, except public holidays and scheduled laboratory closures.



August 25 2005, Jamie Shiers

Annex 3.2. Tier-1 Services

Each Tier1 Centre⁴ forms an integral part of the central data handling service of the LHC 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.

Tier1 services must be provided with excellent reliability, a high level of availability and rapid responsiveness to problems, since the LHC Experiments depend on them in these respects.

The following services shall be provided by each of the Tier1 Centres in respect of the LHC Experiments that they serve, according to policies agreed with these Experiments:

- acceptance of an agreed share of raw and first reconstruction data from the Tier0 Centre, keeping up with data acquisition;
- acceptance of processed and simulated data from other centres of the LHC Computing Grid;
- recording and maintenance of raw, processed and simulated data on permanent mass storage;
- provision of managed disk storage providing permanent and temporary data storage for files and databases;
- provision of access to the stored data by other centres of the LHC Computing Grid and by named analysis facilities as defined in paragraph **Error! Reference source not found.** of this MoU;
- operation of a data-intensive analysis facility;
- provision of other services according to agreed Experiment requirements;
- provision of high-capacity network bandwidth and services for data exchange with the Tier0 Centre, as part of an overall plan agreed amongst the Experiments, Tier1 and Tier0 Centres;
- provision of network bandwidth and services for data exchange with Tier1 and Tier2 Centres, as part of an overall plan agreed amongst the Experiments, Tier1 and Tier2 Centres;
- administration of databases required by Experiments at Tier1 Centres.

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

⁴ The term “Tier1 Centre” includes a distributed Tier1 Centre according to the provisions of this MoU. In terms of services and levels of service a distributed Tier1 Centre shall be indistinguishable from a single-location Tier1 Centre.



August 25 2005, Jamie Shiers

The following parameters define the minimum levels of service:

<i>Service</i>	<i>Maximum delay in responding to operational problems</i>			<i>Average availability⁵ measured on an annual basis</i>	
	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 Centre	12 hours	12 hours	24 hours	99%	n/a
Networking service to the Tier-0 Centre 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 ⁶	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.

Annex 3.3. Tier-2 Services

The following services shall be provided by each of the Tier2 Centres in respect of the LHC Experiments that they serve, according to policies agreed with these Experiments:

- provision of managed disk storage providing permanent and/or temporary data storage for files and databases;
- provision of access to the stored data by other centres of the LHC Computing Grid and by named analysis facilities as defined in paragraph **Error! Reference source not found.** of this MoU;
- operation of an end-user analysis facility;
- provision of other services, such as simulation, according to agreed Experiment requirements;
- provision of network bandwidth and services for data exchange with Tier1 Centres, as part of an overall plan agreed between the Experiments and the Tier1 Centres concerned.

⁵ (time running)/(time scheduled to run)

⁶ Prime service hours for Tier1 Centres: 08:00-18:00 in the time zone of the Tier1 Centre, during the working week of the centre, except public holidays and other scheduled centre closures.



August 25 2005, Jamie Shiers

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:

<i>Service</i>	<i>Maximum delay in responding to operational problems</i>		<i>Average availability⁷ measured on an annual basis</i>
	<i>Prime time</i>	<i>Other periods</i>	
End-user analysis facility	2 hours	72 hours	95%
Other services	12 hours	72 hours	95%

⁷ (time running)/(time scheduled to run)