L.M.Robertson
23 March 2007
Version 7

Summary of the Process for Reporting Experiment Requirements, and Site Capacity and Usage Data for CERN and the Tier-1 Centres

This note summarises recent agreements in the MB on how experiment requests for resources, the capacity installed at CERN and the Tier-1 centres, and the resource usage data are reported by the experiments, the sites and the project.

1. Experiment Resource Requirements

The summary of the experiments' requirements at CERN, in the Tier-1s and in the Tier-2s, formally agreed within the experiments and reported to the C-RRB, is available via a web page linked to the planning page.

The requirements for the current year and the following four years should be specified. The original set of requirements is contained in the computing TDRs of the experiments, and summarised in the LCG TDR. The last update was made in October 2006. These requirements will be updated twice per year:

- in July to allow funding agencies and sites to prepare their resource commitments for the following year, in time for the October meeting of the C-RRB;
- in January, taking account of the experience from the previous year's run.

The requirements in the TDR and later updates to these tables are specified as the *gross* requirements[1]. The *installed capacity* (see section 4) is also specified in gross units and so the requirements and installed capacity can be compared directly.

These requirements are the source of the requirements data used in the Megatable (see section 2) and the Regional Centre Resource Table (see section 3).

2. Megatable

This table specifies relationships between Tier-1 and Tier-2 centres, giving for each experiment inter-site data rates, sizing and utilisation of storage classes, etc. – according to the way in which the experiments expect to use its resources at the site during the 2008 run.

The capacity numbers are all expressed in *gross* units, to be compared with the installed capacity at the site.

[1] After calculating the capacity that the experiment expects to use (the net capacity) the gross requirement is calculated to take account of the efficiency with which the resource can be used in practice in a realistic operational environment, The standard efficiency factors are defined in the TDR as follows, expressed as the usable percentage of the gross capacity: CPU - at Tier-0, Tier-1 and simulation at Tier-2s: 85%; chaotic analysis at the CAF and Tier-2s: 60%; Disk – 70%; Tape – 100%.

The Megatable is maintained by a committee formed by the LCG Resource Coordinator (Chris Eck, chair) and one person appointed by each experiment. Updates are posted on the web labelled with the date of the update, and linked to the Planning page. It is updated regularly. The MB is informed of each new version.

The Megatable contains a section that explains each data field and its provenance.

3. Regional Centre Resource Table (MoU Pledges)

This table, linked to the LCG Planning Page, defines the current status of the resources pledged or planned by funding agencies, organised by regional centre. This includes the current status of Annexes 6.1, 6.2, 6.3 and 6.4 of the MoU.

Changes to MoU pledges are reported by the funding agencies to the LCG Project Office (Chris Eck or Fabienne Baud-Lavigne), after agreeing any substantial changes with the experiments concerned. MoU pledges are expressed in *gross* capacity units.

Each year the Project Office will request new planning figures, but changes may be reported at any time. The pledges for the coming year are formally committed by the funding agencies during the October meeting of the C-RRB (exceptionally the pledges for the 2007 run will be formally committed at the April 2007 meeting).

4. Resource Planning Tables

Requirements

Experiments' medium term requirements for capacity at Tier-1 centres during the current and one to three following quarters are reported through the weekly Experiment Coordination Meeting (ECM) or directly to the Service Coordination Team (Harry Renshall or Jamie Shiers). These requests are expressed as the *gross* capacity that the experiment wishes to use during each quarter – cpu capacity, and storage capacity in each storage class. These numbers are expressed in *gross* terms and so indicate the level of *installed* capacity that the site should make available for the experiment. These numbers should be kept up to date by the experiments.

Installed Capacity and Allocations to Experiments

The installed capacity is reported by sites to the Service Coordination Team during the weekly Operations Meeting or at any other time. Planning targets

should be reported for future quarters, actual capacity for the current quarter. The numbers should be updated whenever they change. The numbers reported should be the *gross* capacity[2] installed, where *installed* is understood to mean tested and configured, ready to be allocated to an experiment at short notice. For storage, in addition, the capacity allocated to each experiment is also specified[3]. Depending on the allocation strategy and policy of each site the sum of the allocations may not equal the installed capacity (e.g. a site may reserve capacity to be allocated according to actual usage or, conversely, a site may adopt an over commitment strategy for space allocation).

The values of the requirements, installed and allocated capacity for the current quarter will be used in the generation of the monthly Accounting Report (see section 5).

The Medium Term Resource Planning Tables are linked to the Service Challenge wiki (→ SC4 Site + Experiment Plans). Current tables: 1Q07 2Q07

5. Resource Usage Data and the Accounting Report

CPU Usage

The CPU Usage and Wall Clock time at Tier-1 sites will be reported by sites to the APEL Repository from the first of March 2007, using units of KSI2K-seconds. The Accounting Report for March and subsequent months will use the numbers extracted from the APEL Repository. No distinction will be made in the report between Grid and Non-grid submitted work.

Disk and Tape Usage

Disk and Tape usage will continue to be reported manually by sites until experience has been gained through the storage accounting pilot and the MB has decided to move to automated reporting of storage usage.

Accounting Report

The monthly accounting report prepared by the LCG project office and linked from the LCG Planning Page. The data for the Accounting Report for CERN and the Tier-1 sites is compiled as follows, using the data present on the first work day on or after the 7th day of the following month.

- The MoU pledge data is taken from the current version of the Regional Centre Resource Table (see section 3).
- The current installed capacity at the site (cpu, disk, tape), and the disk allocation for each VO is taken from the Resource Planning Tables (see section 4).

[2] The gross installed capacity for storage is the total *user accessible space* available – it does not include file system and RAID overheads. For processors the capacity is measured by an agreed benchmarking process.
[3] At a later stage this will be expanded to show allocation by storage class.

- The experiment's requests (cpu, disk, tape) are also taken from the Resource Planning Tables (see section 4).
- The cpu and wall clock usage is taken from the APEL repository (March 2007 onwards).
- The report for the month partially filled as above is sent by the LCG Office (Fabienne Baud-Lavigne) to each site. The site must respond within 10 days providing storage usage data.
- As a temporary measure, until May 2007, the site may also correct data taken from other sources (also reporting the reason for the change).