LHCb Grid usage policy

History

- First draft started: 26 May 2009
- First draft completed: 29 May 2009
- Last update: 3-Feb-10

Introduction

The LHCb Collaboration relies heavily on Grid resources for implementing its Computing Model that has been described in the "LHCb Computing TDR" published in 2005.

The present document further specifies how the resources allocated to LHCb on the Grid should be used. A complementary document specifies the support of the LHCb Computing Project to non-Grid analysis.

The LHCb Computing Model

The LHCb Computing Model as described in the Computing TDR foresees three main types of Grid Computing activities:

- *Real data reconstruction and stripping*:

  After raw data have been recorded at the experimental site, RAW files are transferred to the CERN Tier0 Storage Element (SE) from where they are distributed to the LHCb Tier1 centres[1]. The RAW files are reconstructed at the Tier1s, producing rDSTs that are then used for stripping off events of interest, producing DST files that can be grouped in "streams". DST files are distributed to the LHCb Analysis Centres[2] (LAC). If required, DST files can be further analysed on behalf of Physics Working Groups (PhWG) for producing group-level DST or μDST files.

  This activity is centrally managed by the LHCb Computing Operations team and (μ)DST files are distributed to the LHCb Analysis Centres according to the agreed distribution policy[3].

- *User analysis:*

  Users analyse either real or simulated data from (μ)DST files that have been distributed on the LACs. The output data of user analysis can be stored on Grid storage for further analysis (on the Grid and off-Grid). Grid data analysis runs at any LAC that has access to the required data files. The LAC on which

[1] The LHCb Tier1s are: CNAF, GridKa, IN2P3, NIKHEF, PIC and RAL. Hereafter "Tier1" refers to any LHCb Tier1 or CERN.

[2] The LACs include all Tier1s. Conditions for a non-Tier1 site to be a LAC are described in a separate document.

[3] All (μ)DSTs are distributed to all Tier1s. Distribution to other LACs is under the LAC responsibility.

analysis jobs run is determined by the LHCb Grid system in order to provide the minimum latency. Users don't need to care about the job placement on the Grid. User storage is provided by an SRM space at all Tier1s.

Users may also submit jobs that are not analysing data, e.g. for standalone toy MC studies. Such jobs may run at any Grid site, as no input data is required. Output data if required can also be stored on Grid user storage.

- *Monte-Carlo simulation*

On request from PhWGs, and after approval by the PPG, Monte-Carlo data samples can be simulated, digitised and reconstructed, producing DST files (possibly including MC-truth information). If necessary, stripping can be performed on MC data as for real data. Simulation can be run at any Grid site.

The LHCb Computing Operations team also centrally manages this activity and DST files are distributed to a subset of LACs[4].

Grid access tools

LHCb has developed as Grid interface DIRAC, an integrated Workload (a.k.a. job) and Data Management System. The LHCb Computing project only supports Grid jobs submitted through the DIRAC system. The DIRAC development team is committed to implement all use cases recognised as of general interest for the Collaboration. The LHCb Production System, closely integrated with DIRAC is used for handling all centrally managed activities (a.k.a. production).

User jobs on the Grid should be submitted through the DIRAC System, making use of the ganga[5] tool as front-end user interface. The ganga development team commits to implementing the use cases of general interest to the LHCb Collaboration and to maintain the interface to the DIRAC system as the DIRAC team commits to support and maintain this interface. In addition ganga allows non-Grid job submission for LHCb either on a local machine or to a local batch system. Direct gLite-WMS submission of LHCb jobs is not supported nor allowed.

DIRAC job submission to the Grid uses "generic pilot jobs", i.e. no tasks are submitted to the gLite-WMS with user credentials, but with generic credentials using the VOMS "pilot" role. This paradigm allows LHCb to apply its own policy on the Grid (priorities, users authorisation…). Priorities are assigned by groups of users that are created on request of the PPG that also defines their composition as well as their relative priorities. A user may belong to more than one group and select his group for each type of activity he runs on the Grid.

LHCb Grid sites

All Grid-enabled sites (i.e. properly registered as part of the WLCG production Grid) providing a Computing Element (CE) accepting the LHCb VO can be used for running LHCb Grid jobs (production and user analysis).

---

[4] The Computing Model foresees distribution to CERN and two other Tier1s. Distribution to non-Tier1 LACs is under the LAC responsibility.

[5] ganga is developed in collaboration with ATLAS.

LHCb Grid sites offer resources to the entire Collaboration. These resources are declared to the WLCG management and to the LHCb Computing Resource Management as pledges in favour of LHCb. Resources used on these sites are accounted as such through the standard WLCG accounting system. DIRAC also collects accounting information that is available through the DIRAC Web portal.

The assignment of activities to Grid sites is as follows:

- Simulation production may take place at any Grid site.
- Real data is distributed and reconstructed only at the Tier1 sites.
- Real and simulated data stripping and if needed PhWG-analysis takes place only at the Tier1 sites.
- Analysis not requiring input data may take place at any Grid site.
- User data analysis takes place at LACs only and is data-driven.

Resources accounting

DIRAC provides accounting for the usage of Grid resources: CPU, Storage and Transfers. User accounting is available for CPU and Storage resources.

CPU accounting allows an a-posteriori report on the usage of pledged CPU resources by members of the collaboration. No quota limitations are applied, but it is expected that members of the Collaboration will behave in a collaborative manner and not overload the LHCb Computing system. If necessary the Computing Operations team can give guidance on best practices. Activities requiring large CPU capacity should be handled centrally as productions and therefore agreed as necessary by the PPG and affordable by the OPG.

Storage resources are accounted and a system of disk quota is applied on the overall usage. Users are warned by email when approaching their quota limit and are requested to clean up their storage. If well justified, additional quota can be granted on a case by case basis. The default quota is defined depending on the available resources by the LHCb Data Management and Resource coordination team, and approved by the OPG.

Glossary

- CS: DIRAC Configuration System.
- DIRAC: LHCb Grid interface for Workload and Data management
- Ganga: user interface for job submission, developed jointly by ATLAS and LHCb.
- LAC: LHCb Analysis Centre.
- LFC: LHCb File Catalog.
- LFN: Logical File Name.
- SE: Storage Element on the Grid. It must provide an SRM interface.
- SRM: Storage Resource Manager. A service abstracting storage resources on the Grid.
- WLCG: Worldwide LHC Computing Grid.