

LHCb-PIC agreement for service support

Summary

This document describes the agreement for service support for the following activities of the LHCb collaboration at PIC:

- LHCb DIRAC web site
- LHCb DIRAC Monitoring system
- LHCb DIRAC Accounting system
- LHCb DIRAC Logging system

And this as a long term support agreement for the computing activities of LHCb among both interested parties.

Introduction

DIRAC is the LHCb software infrastructure that manages the use of computing resources for the collaboration. DIRAC is made of a number of collaborating systems, which in term are made of Services and Agents.

From the DIRAC systems, the Workload and Data Management systems that take care of the appropriated scheduling, monitoring and control of the associated tasks run at several servers at CERN, together with the Configuration and Bookkeeping Systems.

For redundancy and error recovery a Request System and part of the Configuration System run on the VO-Boxes at the different LHCb Tier1 sites.

The remaining DIRAC systems (Monitoring and Accounting) have historically been running on servers at PIC due to the important implication of the Spanish groups (Barcelona and Santiago) in the development of the DIRAC project.

Finally, in the last review of the DIRAC project a new system has emerged, the Logging System, which is proposed to run at PIC in the future.

Since the start of the Spanish collaboration to the DIRAC project, there has been full support from PIC to these activities, in term of providing the necessary servers to run DIRAC services. The aim of the present document is to provide a formal agreement to this support that should evolve into long term collaboration among LHCb and PIC with the start-up of the data-taking period.

This document provides an overview of those DIRAC systems that are proposed to be run at PIC in section 1; describes our estimation of their associated requirements in section 2, and finally makes a concrete resource and service agreement for the next couple of years in a final appendix that will have to be reviewed in the future.

1. DIRAC Systems at PIC

In the last years, LHCb has received full support from PIC to run the Accounting and Monitoring Systems in dedicated servers. Also the main DIRAC monitoring web page and one of the two slave Configuration servers have been running at PIC since long.

In the modified DIRAC architecture, after the last project review, Configuration System slave servers are located at the different Tier1 VO-Boxes, providing for

redundancy. A new system has emerged that will allow to keep a certain kind of DIRAC global log, the Logging System. In this scenario LHCb requires support for the following DIRAC systems to be run on dedicated servers at PIC:

- Accounting,
- Monitoring, and
- Logging.

Furthermore, support is also required to run the LHCb DIRAC web site, main entry point for users to access the information provided by the DIRAC systems.

The DIRAC web site

The way to access the information collected by all the DIRAC systems is a number of dedicated web pages which, with appropriated user identification (via grid certificates), provide the way for LHCb and not-LHCb users to query the corresponding system.

All the DIRAC system will allow a certain level of interaction with the user via dedicated web pages. Altogether, these pages form the LHCb DIRAC web site.

For all LHCb users (both individuals, and production or data managers) these pages will be the standard way to monitor the execution of their computing tasks submitted to DIRAC. At the same time it will also provide the external view of LHCb DIRAC setup for external non-LHCb users.

The DIRAC Accounting System

This system is the memory of DIRAC. It provides the means for the different workload and data activities to register their respective use of computing resources according to DIRAC/LHCb view point. The DIRAC components responsible for the execution of the computing tasks register the usage of basic resources, like CPU, Bandwidth..., as well as of more specific ones like number of events processed.

These statistics are kept in a DB backend. Appropriated interface is provided by DIRAC to query this DB in order to produce Accounting reports with the historical view of the usage of computing resources as seen by DIRAC/LHCb.

The DIRAC Monitoring System

This system provides the view of the current activity of the complete DIRAC setup, as well as of each of its systems. Each component provides, at regular intervals, information about a number of the activities it carries out.

This information is collected at a single place, by the Monitoring System that provides means to display it and to obtain different views of the current activities of DIRAC.

The DIRAC Logging System

This new system provides DIRAC with remote syslog functionality. DIRAC components are able to keep important logging messages by sending them to a remote server. The server sorts received messages according to their origin (component and hostname) and level.

The logging messages are kept, for a certain amount of time, in a DB backend and an appropriated interface is foreseen to provide access to accumulated messages.

2. Estimated Requirements for DIRAC Systems at PIC

The expected load on those systems is mainly correlated to the number of tasks (of both Workload and Data Management nature) that DIRAC handles for LHCb. Currently DIRAC manages over 10k jobs per day and a similar order of magnitude of data management operation. This is not an average value, but the current peak value for busy days.

It is expected that this number will grow about 1 order of magnitude once LHC starts operation. Therefore the Accounting System will have to be able to sustain rates up to several Hz of clients registering new entries. Monitoring and Logging Systems are expected to be less demanding provided that in the current design only central servers and agents are sending their data.

The load on the server for the DIRAC web site it is, from our current experience with the site at <http://lhcb.pic.es/DIRAC> not very large at the moment, although as more users join the GRID, the frequency of the queries will certainly increase. Since the web server it is only a mean to interface to the corresponding DIRAC systems running on their own servers it is not expected that it will be a demanding situation.

Given the all the above consideration, LHCb request to PIC, in terms of services that need to be supported, can be summarized as follows:

- One high-performance server for the backend DB support with estimated 1 GB daily amount of new data inserted to the DB and 1-10 Hz of combined insert-select frequency, DB server.
- One medium-performance server for the location of the corresponding DIRAC Servers and Agents associated to the systems running at PIC, DIRAC server.
- One medium-performance servers for the apache server of the DIRAC web site, and local disk shared with the DIRAC server, web server.
- One standard-performance server for development and test of new components and as running hot-spares of either the DIRAC or the web servers.

PIC will be responsible for the operation of the machine, the DB engine and, at least, daily backups for the DB server. For the other servers (DIRAC and web) only the agreed operating system installation (including LCG middleware) will be required from PIC while the rest of the software running on them will be responsibility of LHCb.

Appendix: 2007 concrete service agreement

From the signature of this agreement, and until it is revisited in the future, the requirements described in section 2 correspond to the following hardware and service levels:

- **DB Server:** at least a dual-processor dual-core server providing 8 kSpecInt 2000 of CPU power, 4 GB of system RAM and 500 GB of mirror disk for the DB. It must be possible to extend the available disk space for the DB by a similar amount with minor disruption of the service.
- **DIRAC Server:** at least a dual-processor (or dual-core) providing 4 kSpecInt 2000 of CPU power and 2 GB of system RAM and 150 GB of local disk.
- **Web Server:** at least a dual-processor (or dual-core) providing 4 kSpecInt 2000 of CPU power, 2 GB of system RAM and 150 GB of local disk.

- **Development Server:** at least a dual-processor (or dual-core) providing 3 kSpecInt 2000 of CPU power and 1 GB of system RAM and 150 GB of local disk.

Network connectivity among the servers it is also an important issue, Gigabit connection between either of them is mandatory.

The expected uptime and incidence response for those machines is the same as for LCG Services in the Tier1 infrastructure.