

RESULTS OF THE LHCb EXPERIMENT DATA CHALLENGE 2004

J. Closier, CERN, Geneva, Switzerland
Ph. Charpentier, M. Frank, CERN, Geneva, Switzerland
V. Garonne, A. Tsaregorodtsev, CPPM, Marseille, France
M. Witek, Henryk Niewodniczanski Institute of Nuclear Physics, Cracow, Poland
V. Romanovski, IFVE, Protvino, Russia
U. Egede, Imperial College, London, United Kingdom
V. Vagnoni, INFN, Bologna, Italy
I. Korolko, ITEP, Moscow, Russia
J. Blouw, MPI, Heidelberg, Germany
G. Kuznetsov, G. Patrick, RAL, Didcot, United Kingdom
M. Gandelman, UFRJ, Rio de Janeiro, Brazil
R. Graciani-Diaz, Universidad de Barcelona, Barcelona, Spain
R. Bernet, Universität Zürich, Zürich, Switzerland
N. Brook, University of Bristol, Bristol, United Kingdom
A. Pickford, University of Glasgow, Glasgow, Scotland
M. Tobin, University of Liverpool, Liverpool, United Kingdom
A. Saroka, I. Stokes-Rees*, University of Oxford, Oxford, United Kingdom
J. Saborido-Silva, M. Sanchez-Garcia, USC, Santiago de Compostela, Spain

Abstract

The LHCb [1] experiment performed its latest Data Challenge during 2004; the first phase (Monte Carlo production) lasted from May to August 2004.

The main goal of the Data Challenge was to stress test the LHCb production system and to perform distributed analysis of the simulated data.

The LHCb production system called DIRAC [2] provided all of the necessary services for the Data Challenge: Production and Bookkeeping databases, File catalogs, Workload and Data Management systems, Monitoring and Accounting tools. It enabled resources of more than 20 LHCb production sites as well as the LCG2 grid resources to be accessed in a consistent way. 187 Million events constituting 62 TB of data were produced and stored in 5 Tier 1 centers.

This paper will present the Data Challenge results, the experience gained utilising DIRAC native sites and the LCG grid as well as further developments necessary to scale to the ultimate needs of the LHCb experiment.

DATA CHALLENGE 2004 GOALS

Computing goals

The main goal of the LHCb Data Challenge 04 is to gather information which will be used to validate the

* Marie Curie Training Site Fellow at CPPM, Marseille, France.

LHCb distributed computing model, including distributed analysis. Software as realistic as possible in terms of performance was used to test the robustness of the LHCb applications and the production system. The LCG application software was used for the first time in the Data Challenge. The goal of LHCb was to use the LCG resources to provide at least 50% of the total production capacity.

The Data Challenge was divided into three phases:

- Production : Monte Carlo simulation
- Stripping : Event pre-selection
- Analysis

Physics goals

To test a realistic analysis environment, the data produced will be used for physics & High Level Trigger studies. These studies require 30×10^6 signal events, 15×10^6 events of specific backgrounds and 125×10^6 background events where the ratio between B inclusive and minimum bias is 1:1.8. This represents a factor 3 increase with respect to the number of events produced during the Data Challenge 2003.

PRODUCTION MODEL

Model description

The production for the Data Challenge 04 used the Distributed Infrastructure with Remote Agent Control (DIRAC) [2] system. DIRAC was used to control

resources both at DIRAC dedicated sites and those available within the LCG environment.

A number of central services were deployed to serve the Data Challenge. The key services are:

- A production database where all prepared jobs to be run are stored
- A Workload Management System [3] which dispatches jobs to all the sites according to a “pull” paradigm
- Monitoring and accounting services [4] which are necessary to follow the progress of the Data Challenge and allow the breakdown of resources used
- A Bookkeeping service and the AliEn [5] File Catalog keep track of all datasets produced during the Data Challenge.

The technologies used in this production are based on C++ (LHCb software), python (DIRAC tools) and XML-RPC (the protocol used to communicate between jobs and central services). ORACLE and MySQL are the two databases behind all of the services. Oracle used for the production and bookkeeping databases [6], and MySQL for the workload management system.

Strategy

LHCb distinguishes two kinds of resource:

- DIRAC native sites where computing resources are dedicated to LHCb
- LCG site where LHCb is allowed to use some share of LCG computing resources

DIRAC is deployed on each site which is participating in the Data Challenge. On DIRAC native sites, the DIRAC tarball is installed manually and the site manager configured the production system. Any site, where DIRAC has been deployed, is allowed to participate in the Data Challenge once it has been tested successfully.

On the LCG, “agent installation” jobs were submitted continuously. These jobs check if the Worker Node where they end up is eligible to run an LHCb job. If the assessment is positive, the job installs the DIRAC agent, which then executes as on any other DIRAC site within the time limit allowed for the job. This mode of operation on LCG allowed the deployment of the DIRAC infrastructure on LCG resources and using them together with other LHCb Data Challenge resource in a consistent way.

During the production phase, we attempted to always keep jobs in the queue. To achieve this, each site is encouraged to run continuously the local DIRAC agent via cron jobs or *runsv* [7] or daemon. For LCG, “agent installation” jobs were submitted continuously via a cron job on the LCG User Interface.

DATA STORAGE

A typical LHCb job is composed of three steps which are:

- Simulation which produced SIM files
- Digitisation which produced DIGI files
- Reconstruction which produced DST files

Of the data produced, only that from the reconstructed step (DSTs) was kept and distributed to the Tier-0 at CERN and the “nearest” Tier-1’s. Tables 1 and 2 show the distribution of DSTs among the resource centres.

Table 1: Tier0 storage

Tier0	Nb of events	Size (TB)
CERN	187 557 231	62

Table 2: Tier1 storage

Tier1	Nb of events	Size (TB)
CNAF (Italy)	37 129 350	12.6
RAL (UK)	19 462 850	6.5
PIC (Spain)	16 505 010	5.4
FZK (Germany)	12 486 300	4
Lyon (France)	4 368 656	1.5

PERFORMANCE

The production phase of the LHCb Data Challenge is completed; 187 million events were produced. The Data Challenge commenced with only DIRAC native sites and after two weeks LCG resources were gradually incorporated. During this period, around 1.8 million events per day were produced. LCG production was paused and an interim report on our experience was delivered to the LCG Grid deployment team.

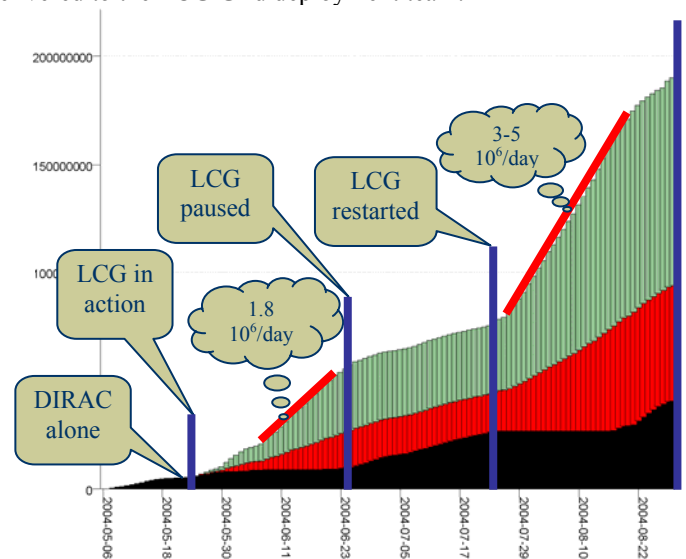


Figure 1: Number of events integrated during the production phase.

The grid deployment team addressed many of the issues raised by LHCb during the pause and the

improvements were evident after the resumption of LCG use, as is illustrated in Figure 1.

After the pause, the production was around 3 to 4 million per day with a peak to 5 million per day as shown in Figure 2. LHCb were able to run up to 3500 jobs concurrently




Figure 2: Daily performance.

During the production phase, 20 DIRAC native sites and more than 43 LCG sites provided computing resources; amongst these 43 LCG sites, 8 were also DIRAC sites. It should be noted that resources from non-LHCb countries not participating in LHCb have been used inside LCG, which produced more than 50 million events.

At the end of the production phase of the Data Challenge, half of the number of events has been produced with LCG resources. In figure 3, it can be seen that in August more than 75% of the events were produced with LCG resources.

About 15% of the production on LCG was performed on relatively small sites, which demonstrates adaptability of the LHCb approach to the utilisation of the LCG resources.




Figure 3: DIRAC – LCG share.

In total, 211000 jobs were submitted to LCG. LHCb cancelled 26000 jobs because the lifetime of the proxy could expire before the end of a job which has queued too

long in the LCG Resource Broker. Details of the problems encountered are given in Table 3.

Table 3: LCG efficiency

	Jobs (k)	%Remain
Submitted	211	
Cancelled	26	
Remaining	185	100.0%
Aborted (not run)	37	20.1%
Running	148	79.7%
Aborted (run)	34	18.5%
Done	113	61.2%
Retrieved	113	61.2%

A large fraction of jobs in status “Aborted” were due to site misconfiguration before LHCb jobs started to run (20%). However even after jobs started running, a large fraction were aborted (18.5%), some due to the fact that the published information about CPU time limit was not normalized. From the 113000 jobs which reached the status “Done”, the analysis of the output sandbox show that 69000 jobs finished successfully. In fact there were problems where the incorrect status in LCG was reported; the LHCb bookkeeping reported 81000 LCG successful jobs registering data.

Table 4 shows the status of a job from the user’s point of view when they have the status “Done” for LCG. From the 44000 jobs which did not finish successfully, the LHCb application errors represent only 1.85% (including possible site configuration problems that caused an application crash e.g. exceeding a disk quota) while 14.86% are due to initialization error (e.g. missing python, failing DIRAC installation, failing software installation or failing connection to DIRAC servers).

Table 4: output sandbox analysis of jobs in status “Done” for LCG

	Jobs (k)	% Retrieved
Retrieved	113	100.00%
Initialization error	17	14.86%
No Job in DIRAC	15	13.06%
Application error	2	1.85%
Other Error	10	9.00%
Success	69	61.23%
Transfer Error	2	1.84%
Registration Error	1	0.64%

The LHCb Data Challenge shows that the LCG efficiency was around 61%.

LESSONS

This Data Challenge demonstrates that the concept of light, customizable and easy to deploy DIRAC agents is very effective. Once the agent is installed, it can effectively run as an autonomous operation and it needs only one single configuration file. The procedure to update or to propagate bug fixes for the DIRAC tools is quick and easy but care is necessary to ensure the

compatibility between DIRAC releases and ongoing operations.

To distribute the LHCb software, the “paratrooper” approach has been chosen, which means that the installation of the software is triggered by a running job and the distribution contains all the binaries and is independent of the Linux flavour. Nevertheless, new services to keep track of available and obsolete packages and a tool to remove software packages would be desirable.

The DIRAC system relies on a set of central services. Most of these services were running on the same machine that ended up with a high load and too many processes. With thousands of concurrent jobs running in normal operation, the services are approaching a Denial of Service regime, where the response slows down and services become stalled.

In future releases of the DIRAC system, the approach to error handling and reporting of failures of different services should be improved.

As LCG resources were used for the first time, several areas have been identified where improvements should be made.

- Many jobs failed or aborted because the normalisation of CPU time was not homogeneous across LCG.
- It was impossible to do bulk operations with the LCG commands.
- The mechanism for uploading or retrieving OutputSandbox should be improved, in particular to have information about jobs with status “Failed” or “Aborted”.
- The collection of log file information and the navigation through the log files should be improved to ease debugging of problem.
- The management of each site should be reviewed to detect and to take action when a misconfigured site becomes a “black-hole”.
- The publication of information about site intervention should be also provided to the Resource Broker or to the Computing Element.

CONCLUSIONS

The LHCb Data Challenge production phase is completed. The production target was achieved with 187 million events in 376 kSI2k years on which 50% was produced with LCG resources. The LHCb strategy of submitting “empty” DIRAC agents to LCG has proven to be very flexible allowing a success rate above LCG alone. Areas for improvement in both DIRAC and LCG have been identified. In particular, DIRAC needs to improve the reliability, and associated redundancy, of the servers and LCG needs improvement on the single job efficiency. In both cases, extra protection against external failures, e.g. network or unexpected shutdowns, must be built in.

ACKNOWLEDGEMENTS

The success of the production phase is due to the dedicated support from the LCG deployment team and the DIRAC site managers.

REFERENCES

- [1] LHCb, S. Amato et al, LHCb Technical Design Report, CERN-LHCC-2003-030, September 2003. <http://eu-datagrid.web.cern.ch/eu-datagrid>
- [2] <http://dirac.cern.ch/>. A. Tsaregorodtsev et al, “DIRAC - The Distributed MC Production and Analysis for LHCb”, CHEP’04, ID 377
- [3] V. Garonne et al, “DIRAC Workload Management System”, CHEP’04, ID 365
- [4] M. Sanchez Garcia et al, “A Lightweight Monitoring and Accounting System for LHCb DC04 Production”, CHEP’04, ID 388
- [5] AliEn, <http://alien.cern.ch>
- [6] C. Cioffi et al, “File-Metadata Management System for the LHCb Experiment”, CHEP’04, ID 392
- [7] G. Pape, Runit Service Supervision Toolkit, <http://smarden.org/runit>