FERMILAB-CONF-11-866-CD

Large scale and low latency analysis facilities for the CMS experiment: development and operational aspects

H Riahi[1], S Gowdy[3], P Kreuzer[2], J Bakken [4], M Cinquilli[3], D Evans[4], S Foulkes[4], R Kaselis[3], S Metson[5], D Spiga[3], E Vaandering[4]

1 Università and INFN Perugia, 2 Rheinisch-Westfaelische Tech. Hoch. (RWTH), 3 CERN, 4 FNAL, 5 University of Bristol

E-mail: `Hassen.Riahi@pg.infn.it`

Abstract.

While a majority of CMS data analysis activities rely on the distributed computing infrastructure on the WLCG Grid, dedicated local computing facilities have been deployed to address particular requirements in terms of latency and scale.

The CMS CERN Analysis Facility (CAF) was primarily designed to host a large variety of latency-critical workflows. These break down into alignment and calibration, detector commissioning and diagnosis, and high-interest physics analysis requiring fast turnaround. In order to reach the goal for fast turnaround tasks, the Workload Management group has designed a CRABServer based system to fit with two main needs: to provide a simple, familiar interface to the user (as used in the CRAB Analysis Tool[7]) and to allow an easy transition to the Tier-0 system. While the CRABServer component had been initially designed for Grid analysis by CMS end-users, with a few modifications it turned out to be also a very powerful service to manage and monitor local submissions on the CAF. Transition to Tier-0 has been guaranteed by the usage of the WMCore, a library developed by CMS to be the common core of workload management tools, for handing data driven workflow dependencies. This system is now being used with the first use cases, and important experience is being acquired. In addition to the CERN CAF facility, FNAL has CMS dedicated analysis resources at the FNAL LHC Physics Center (LPC). In the first few years of data collection FNAL has been able to accept a large fraction of CMS data. The remote centre is not well suited for the extremely low latency work expected of the CAF, but the presence of substantial analysis resources, a large resident community, and a large fraction of the data make the LPC a strong facility for resource intensive analysis.

We present the building, commissioning and operation of these dedicated analysis facilities in the first year of LHC collisions; we also present the specific development to our software needed to allow for the use of these computing facilities in the special use cases of fast turnaround analyses.

1. Introduction

The central goal of the Compact Muon Solenoid (CMS) experiment[1] is to explore physics at the TeV energy scale, exploiting the collisions delivered by the Large Hadron Collider (LHC)[2] at CERN. CMS utilises a tiered and distributed infrastructure of computing centres to perform analysis on the collected data [3].

Published under licence by IOP Publishing Ltd

Operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy

The Tier-0 centre at CERN accepts data from the CMS online system, archives the data and performs prompt first pass reconstruction. The time-critical analysis workflows, in particular calibration and alignment, which result in improved knowledge of detector, are executed always at CERN using the dedicated Analysis Facility (CAF). Reconstructed data at the Tier-0 together with the corresponding raw data are distributed to Tier-1s which take custodial responsibility. At the Tier-1 centres data is re-processed based on improved detector conditions and selected on the interesting physics channels. The analysis datasets resulting are then sent to Tier-2s for end user analysis.

While a majority of CMS data analysis activities rely on the distributed computing infrastructure on the Worldwide LHC Computing Grid (WLCG)[4], dedicated local computing facilities have been deployed to address particular requirements in terms of latency and scale.

The CAF plays a unique role for latency-critical functions at the source of the analysis chain, with no dependence on the Grid infrastructure. Such workflows are based on a highly selected subset of the full production data, so called Calibration and Express Streams (ES), and are made available to restricted analysis groups soon after data taking occurred. They can be classified as follows in order of priority:

- alignment and calibration;
- trigger and detector diagnostics, monitoring and performance analysis;
- physics monitoring, analysis of express stream, fast-turnaround high-priority analysis.

The key hardware components for addressing the low-latency and high-priority requirements are a large dedicated processing farm and a highly accessible disk-only storage system. The limitation in available resources demands selectiveness in the jobs to be run and in the data to be stored.

In addition to the CERN CAF facility, FNAL has CMS dedicated analysis resources at the FNAL LHC Physics Center (LPC), LPC Central Analysis Facility (LPC CAF). In the first few years of data taking FNAL will be able to accept a large fraction of the CMS data. The remote centre is not well suited for the extremely low latency work expected of the CAF, but the presence of substantial analysis resources, a large resident of CMS community, and a large fraction of the data make the LPC a strong facility for resource intensive analysis. The key hardware components for handling the huge number of CMS users accessing a large fraction of data stored at the FNAL Tier-1 for analysis is a large local computing farm dedicated to CMS user analysis.

Executing the offline workflows for alignment and calibration represents a complex and time consuming task, with intricate data dependencies and stringent latency requirements. In addition, there is a need to reduce the delays executing the typical analysis workflows in order to be able to run analysis jobs as soon as data are available at the various computing sites. These are the main motivations which led CMS to develop a fully automated system for managing analysis and calibration workflows. This system is a tool designed to feed information to components that take care of the distributed data bookkeeping, and then triggering the submission of jobs once user conditions are satisfied. This system is now being used with the first use cases and important experience is being acquired.

In this paper we review the resources in the most important prompt analysis facilities for CMS, namely CERN CAF and FNAL LPC CAF, followed by respective performance in the first year of the LHC start-up. We also briefly present the motivations and the design of the automation system. We describe the experience supporting the first use case, the calibration of the Beam Spot measurement[5], as well as the efficiency of the automation system.

2. CERN Analysis Facility

2.1. CAF resources

The technical requirements on the CAF are based on the capability to satisfy high-priority and low latency workflows, by means of enforcing controlled and prioritised access policies to hundreds of users. The CAF systems need to be setup such that particular jobs can start with very short latency in parallel to many other concurrent jobs and with highly efficient access to the CMS production data.

A dedicated processing farm has been installed in various steps since 2008, reaching a total of 876 cores. Access, scheduling and sharing is handled by the Platform LSF[17] system. The batch system comprises a set of queues (788 job slots) available to all CAF users and a special queue (54 reserved job slots) dedicated to very high-priority alignment and calibration workflows.

A large disk storage system based on CASTOR technology[6] has been installed in various steps since 2008, reaching a capacity of 2.1 PB. Highly efficient and asynchronous access to the data coming from the Tier-0 is ensured by the disk-only characteristic of this storage pool, hence avoiding any kind of tape staging latency, which is a typical bottle neck encountered for large storage systems. In addition, CAF sub-groups and users are sharing a 9 TB AFS space for POSIX-like storage needed for analysis. Finally, a set of 5 worker nodes (40 cores, 50 slots) are reserved for interactive access to CAF batch and storage resources, and two 8-core computers are hosting the CRAB Servers for job submissions.

The CAF user management is handled via the "LSF WEB" application provided by CERN/IT: the tool is used for user registration and to set the fair share among CAF subgroups. It is also automatically interfaced to the main CAF resources (CASTOR, AFS, CRAB Server, Email lists) for granting access permission to CAF users. This tool turned out to be very convenient for user management, since it represents the single entry point to any CAF resource.

2.2. CAF performance

The main CAF data management and workflow management components include the job submission tool "CMS Remote Analysis Builder" (CRAB)[7], the Data Bookkeeping System (DBS)[8] and the Data Transfer system (PhEDEx)[9]. These are standard CMS tools, with only minor adaptations compared to those used for distributed analysis on the GRID.

The CAF batch queues rely on a fair share system, in which ∼30 CAF subgroups are allocated resources averaged over time. The total number of jobs submitted during a six months period after the beginning of LHC collisions data taking was $\sim 2\text{x}10^6$, while the cumulative number of active users has reached nearly 300 during the same period, see Figure 1. The main CPU management characteristics are a split into 3 batch queues (with maximum CPU time of 1 hour, 1 day and 1 week, respectively) and a limit of maximum 100 running jobs/user for the two shorter queues, to avoid batch congestions by a single user, while this limit is reduced to 10 for the longer queue.

Many CAF users have been running their workflows with a standalone version of CRAB in recent years and hence the migration of their workflows to the CRAB Server application is straightforward. In the same spirit, migration of CAF workflows from CERN to the GRID is very simple, since based on a quasi identical CRAB configuration; this is of particular importance for the prioritization of CMS analysis workflows and for an optimal usage of resources between CERN and the distributed GRID infrastructure. In 2010, the main CRAB Server applications running on the CAF were related the beam spot calculation and the Strip Tracker calibration.

No major bottle-neck was encountered with the CAF batch system during the first year of LHC data taken, while occasional issues were observed:

- Low Job CPU efficiency (<50%), limiting the overall CPU utilisation performance.
- Very large amount of submissions at a given time by single users, sometimes affecting the latency needs of concurrent workflows

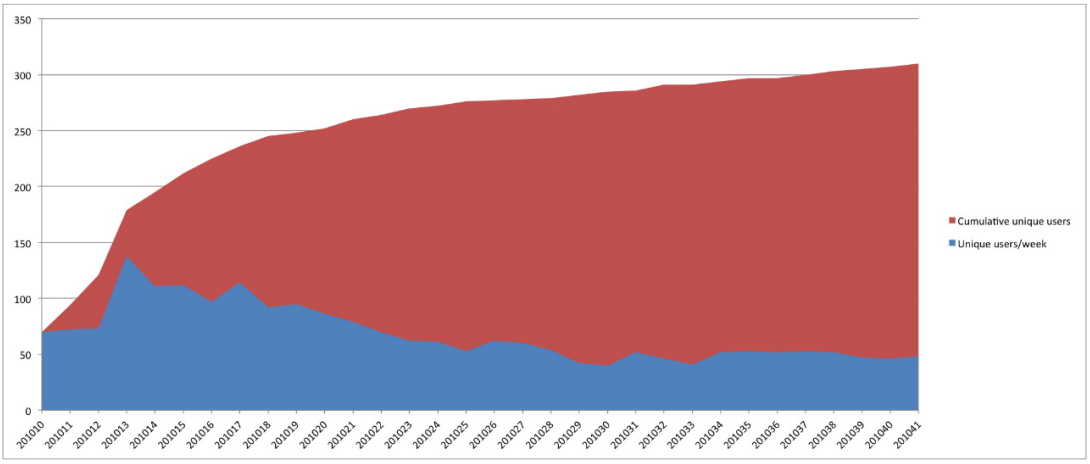

Figure 1. Number of distinct and cumulative CAF users per week, during the first 6 months LHC collisions data taking.

- Occasional disk server failures

It is suspected that the low CPU efficiency is partially due to I/O and this can be addressed switching from rfio[6] to xrootd[10] as the protocol to access CASTOR data.

The latency to get the express data onto CAF has been satisfactory (average about half an hour) and this is described in another submission to these proceedings[11].

3. FNAL LPC CAF

3.1. LPC CAF resources

The key design of the LPC CAF is that users can transparently read the data stored at the FNAL Tier-1. Given the large community residing at FNAL and the massive amount of data stored at the FNAL Tier-1, the technical requirement on the LPC CAF is based on the capability to handle huge number of local users reading the large fraction of data stored at FNAL Tier-1 by making available a large and dedicated computing facility, not Grid accessible.

A processing farm has been installed in several steps, reaching a total of 355 worker nodes and 2484 cores. Access, scheduling and sharing is handled by the Condor[18] scheduler system. The LPC CAF worker nodes can receive 1 job per core, leading to 2484 available job slots.

The large FNAL Tier-1 storage system is based on dCache[12] technology. Each LPC CAF user and group has a quota controlled area in this system. Moreover, the LPC CAF shares with the Tier-1 ∼300 TB of filesystem based on BlueArc[13] technology. The later is used to host user home area and data area.

3.2. LPC CAF performance

For the last 6 months, the cumulative number of active LPC CAF users so far has reached ∼570, which represents ∼20% of the whole CMS community, and more than 1 PB of CMS data has been transferred to FNAL Tier-1. In general, the data is available at FNAL within a short latency (order of hours/days) after the data is taken. In this same period, more than 2,500,000 users analysis jobs have exploited this facility. Such contribution makes the LPC CAF a strong facility in the CMS computing infrastructure for resource intensive analysis. This is illustrated in Figure 2.

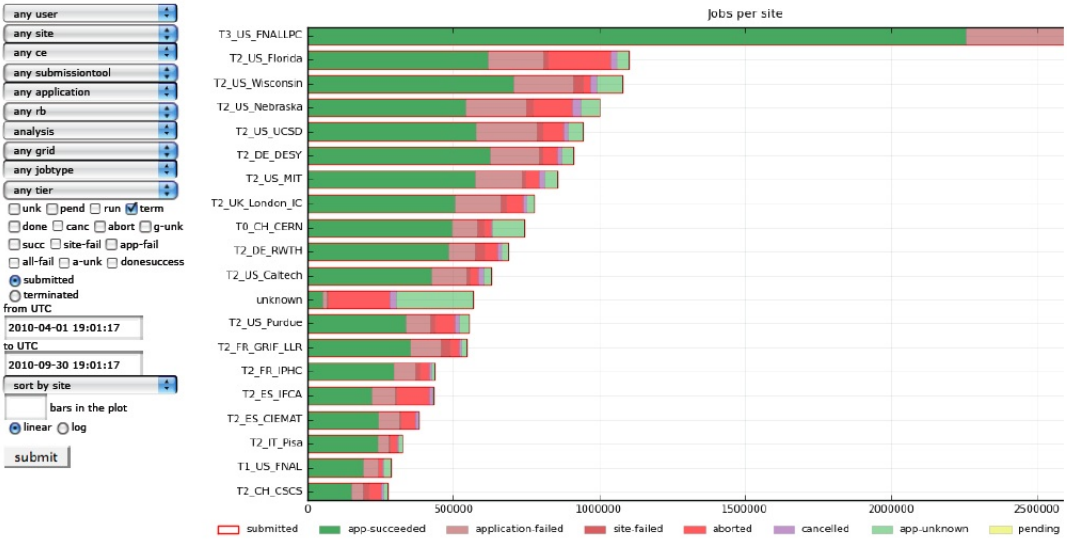

Figure 2. FNAL LPC contribution in CMS computing infrastructure

4. CRABServer based automation system

4.1. Requirements for automated processing of workflows

The goal of this work is to provide an automatic way of managing CMS data analysis workflows. The basic concept of the automation is that workflows should be executed in a data driven manner. It foresees to take actions automatically when a sufficient amount of suitable new data are ready. The actions typically involve submitting jobs as soon as sufficient new data are accumulated, which can be defined as a number of new files or new luminosity sections in a run, the end of a run, or as a given number of events in a file etc.

The foremost objective is to automate alignment and calibration workflows for the following reasons:

- during the initial data-taking, the workflows are not mature enough to be handed over to a central data operations team;
- the alignment and calibration team need to be able to monitor the execution of their workflows very closely and apply correctness to the methods themselves and their configurations;
- and with less priority, to reduce the delays supported by typical analysis workflows.

The general aim is to automate the work done on behalf of users reducing the efforts that they invest to manage their analysis workflows and decreasing, of course, the delays in their execution.

4.2. Implementation strategy

The architecture of the new CMS Workload Management system is based on the new WMCore/WMAgent frameworks[14] and makes use of existing CMS tools, in particular CRAB.

WMCore/WMAgent is the common workload management platform, including core libraries, for managing jobs in CMS. It is designed to support three main application areas: data processing and reconstruction, large scale production of Monte Carlo simulated events and data analysis. CRAB is a user-friendly interface, built on a client-server model, to manage data analysis jobs in a local or distributed environment hiding the complexity of interactions with the Grid and

CMS services. In the future, it is intended that CRAB itself will be implemented in these new common workload management components.

In order to build the infrastructure needed to support automated workflows, an interface has been developed in the current CRABServer framework to interact with WMAgent using WMBS (Workload Management Bookkeeping system).

The key feature of the design strategy is to exploit the WMBS system in CRABServer. WMBS is a library provided as part of WMCore framework, which is by design fully usable as part of the Tier-0 architecture. This approach is exploited to realise the migration of alignment and calibration workflows into the Tier-0 system. The use of CRAB allows us to benefit from its familiar interface to computing resources for physicists and also to achieve the first step required to build the next generation of the CRAB framework (based on WMCore/WMAgent).

4.3. Experience supporting the first workflow and results

The general calibration strategy for CMS is the migration from "standard" workflows running on the CAF to "prompt calibration" workflows running on the Tier-0. The first workflow that is being supported by the automated system is the beam spot determination. The beam spot is the luminous region produced by the collisions of the LHC proton beams. Knowing the position of this is important for accurate physics measurements, and it needs to be measured precisely in an automated fashion for correct offline data reconstruction.

For the beam spot workflow, the data source is the Tier-0. The T0AST Feeder is developed in order to regularly poll the Tier-0 Data Aggregation System (DAS)[15] interface for a given workflow. T0AST Feeder injects new files once they appear in the Tier-0 into the WMBS instance.

Measuring the beam spot using CRABServer required the development of four feeders:

- T0ASTChain allows to submit jobs over outputs of another workflow as soon as they appear in the CAF storage system. This feeder allows users to iterate automatically between their workflows;
- T0ASTRun allows to submit jobs over runs as soon as they are considered as finished by the Tier-0 system;
- T0AST allows to submit jobs over files as soon as they appear in the CAF storage system;
- DBSFeeder allows to submit jobs over files as soon as they appear in DBS. It has been developed to automate also the other kinds of alignment and calibration workflows (Strip Tracker calibration workflows...), and the typical analysis workflow. It is the feeder used for Grid jobs.

All WMBS Feeders inherit from the super-class Feeder. The Feeder based architecture allows to add a new feeder whenever needed by just writing a class which inherits from Feeder parent class and using correctly the WMBS library. Figure 3 shows the general design, using the Unified Modeling Language (UML)[16], of the FeederManager module.

Since the support of alignment and calibration workflows represents a high priority task for the CMS collaboration, the T0Emulator daemon has been developed in order to reduce the timescales to reach this goal. It allows running the machinery during the down periods of the Tier-0, in order to develop and test the beam spot workflow during these LHC down times.

By the 2010 LHC start up, a dedicated production instance of CRABServer running the automation machinery has been setup to serve the alignment and calibration workflows.

Given that the alignment and calibration workflows are running at CAF and thus were limited limited by the Andrew File System (AFS) token lifetime, the immediate feedback was asking to extend the lifetime of the workflows to be as long as possible. The ability to initialise a renewable token for 3 days was added in CRAB client and the automatic user token renew was included

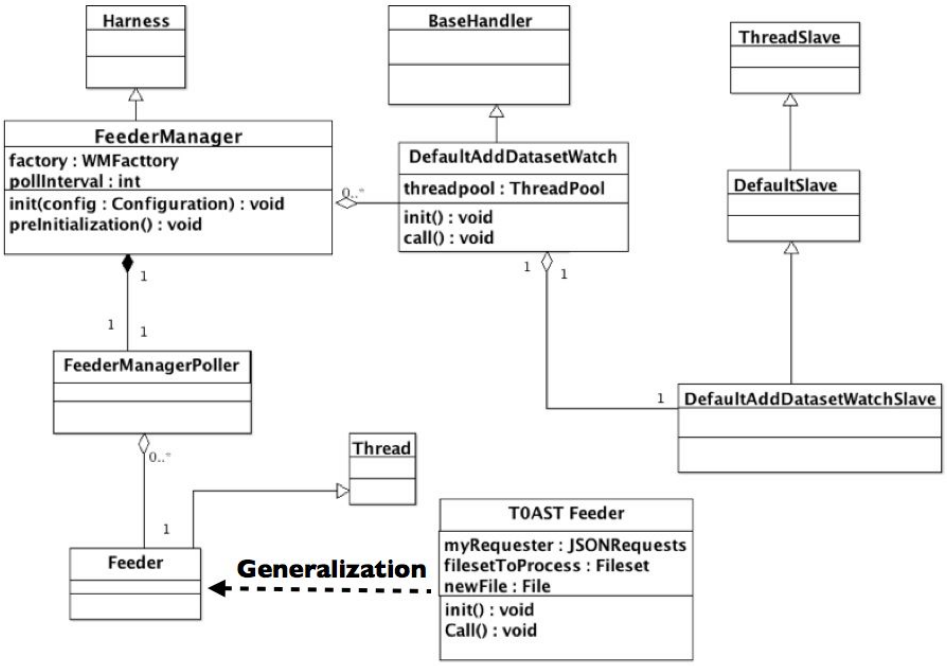

Figure 3. General Class diagram of the FeederManager

in CRABServer. The renew period was extended in a second step to one month (maximum allowed at CERN).

Monitoring the calibration of the beam spot workflow when the maximum user token validity is 3 days, the machinery takes less than 2 hours between the start of collision runs processing in Tier-0 and the beginning of the automatic jobs submission in CRABServer over the data that is produced by these runs. Users submit their workflows when they know that there will be collisions in the next few days. Figure 4 shows this time interval for runs from 138560 to 140317. Such latency allows alignment and calibration experts to update the conditions database with beam spot constants on the time scale required by CMS. Data from non collision runs are not used for the beam spot calculation. In the future the feeder will be improved to filter out data taken during non collision runs.

5. Conclusion

Nearly 1000 users were performing prompt analysis and latency critical activities at various sites during the last 6 months of the first year of LHC data taking.

The CMS CAF has been successfully used during this first year of LHC data taking. The highly accessible disk-only storage system and the flexible batch system have met the low latency and high priority requirements, without encountering any major bottle neck.

Supporting the largest fraction of analysis jobs, LPC CAF farm is playing a key role in CMS computing infrastructure

The CRABServer based automation system has been used with success to automate alignment and calibrations workflows at CAF and has shown satisfactory performance. By automating

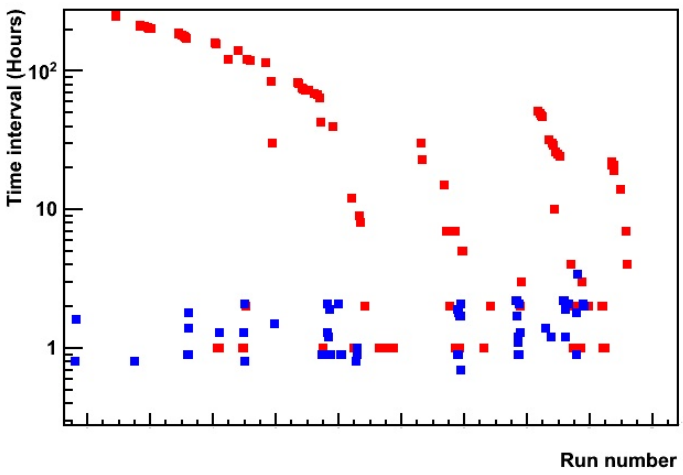

Figure 4. Time interval between Tier-0 processing and jobs submission in CRABServer from run 138560 to 140317.
■ non collision runs
■ collision runs

workflows and standardising jobs, this tool will help to increase the CPU efficiency on the CAF.

Within CRABServer/WMBS, it has provided a clear migration route for alignment and calibration workflows into the more structured and controlled Tier-0 system [11].

Automated workflows will be naturally supported by the new frameworks, WMCore/WMAgent. The fact that CMS is moving to the next generation of data management tools can lead to automation support using other sites such as the FNAL LPC which has proven to be a very reliable environment for analysis. This will lead to a large number of users efficiently performing prompt and latency critical analyses.

References

[1] The CMS Collaboration, Adolphi R *et al* 2008 The CMS experiment at the CERN LHC *JINST* **3** S08004
[2] LHC Homepage http://cern.ch/lhc-new-homepage/
[3] CMS Collaboration 2004 *The CMS Computing Model* CERN/LHCC 2004-035
[4] *LHC Computing Grid (LCG)* Technical Design Report LCG-TDR-001 CERN/LHCC 2005-024, (2005)
[5] Spalding J, Collaboration for two CMS working groups: the Machine Interface Group and Trigger Simulation Group 2008 *Background Issues for CMS* LHC Workshop on Experimental Conditions and Beam-Induced Detectors Backgrounds (CERN, Geneva, Switzerland) pp.60-64
[6] Presti L *et al* 2007 *CASTOR: A Distributed Storage Resource Facility for High Performance Data Processing at CERN* 24th IEEE Conf. on Mass Storage Systems and Technologies
[7] Spiga D *et al* 2007 *The CMS Remote Analysis Builder (CRAB)* 14th Int. Conf. on High Performance Computing (Goa, India) vol. 4873 pp. 580-586 ISBN/ISSN: 978-3-540-77219-4
[8] Afaq A *et al* 2007 The CMS dataset bookkeeping service *J.Phys.Conf.Ser* **119** 072001
[9] Tuura L *et al* 2008 Scaling CMS data transfer system for LHC start-up *J.Phys.Conf.Ser* **119** 072030
[10] Moneta L *et al* 2008 *Recent developments of the ROOT mathematical and statistical software J.Phys.Conf.Ser* **119** (2008) 042023
[11] Hufnagel D *et al* 2010 *The architecture and operation of the CMS Tier-0* in these proceedings
[12] dCache web site: http://www.dcache.org
[13] BlueArc web site http://www.bluearc.com/
[14] Wakefield S *et al* 2008 *Large Scale Job Management and Experience in Recent Data Challenges within the LHC CMS experiment* Proceedings of XII Advanced Computing and Analysis Techniques in Physics Research (Erice, Italy)
[15] Kuznetsov V *et al* 2010 *The CMS Data Aggregation System* Procedia Computer Science **1** 1529-1537
[16] Larman C 2005 *UML 2 et les Design Patterns* Pearson Education ISBN: 2-7440-7090-4
[17] LSF web site http://www.platform.com/Products/platform-lsf
[18] Condor web site http://www.cs.wisc.edu/condor/