FERMILAB-CONF-00-068-CD

Experience in testing the Grid based Workload Management System of a LHC experiment

V. Miccio[a b c], A. Fanfani[c d], D. Spiga[e b], M. Cinquilli[e], G. Codispoti[d], F. Fanzago[b c], F. Farina[f b], S. Lacaprara[g], E. Vaandering[h], A. Sciabà[b] S. Belforte[i]

[a]**Corresponding/Contact Author**
CERN, BAT.28-1-019, 1211 Geneve 23, vincenzo.miccio@cern.ch Office: +41 (0)22 76 77215 − Fax: +41 (0)22 76 79330

[b]CERN, Geneva, Switzerland

[c]INFN/CNAF, Bologna, Italy

[d]University Bologna, Italy

[e]University and INFN, Perugia

[f]INFN, Milano-Bicocca, Italy

[g]INFN Legnaro,Italy

[h]FNAL, Batavia, Illinois, USA

[i]INFN, Trieste, Italy

The computational problem of a large-scale distributed collaborative scientific simulation and analysis of experiments is one of the many challenges represented by the construction of the Large Hadron Collider *(LHC) at the* European Laboratory for Particle Physics *(CERN). The main motivation for LHC to use the Grid is that CERN alone can supply only part of the needed computational resources, while they have to be provided by tens of institutions and sites. The key issue of coordinating and integrating such spread resources leads the building of the largest computing Grid on the planet.*

Within such a complex infrastructure, testing activities represent one of the major critical factor for deploying the whole system. Here we will focus on the computing model of the Compact Muon Solenoid *(CMS) experiment, one of the four main experiments that will run on the LHC, and will give an account of our testing experience for what concerns the analysis job workflow.*

Keywords: Grid Computing, Distributed Computing, Grid Application Deployment, High Energy Physics Computing

Conference: GCA'08 − The 2008 International Conference on Grid Computing and Applications

1. Introduction

CMS represents one of the four particle physics experiments that will collect data at LHC starting in 2008 at CERN, and one of the two largest collaborations. The outstanding amount of produced data − something like 2 PB per year − should be available for analysis to world-wide distributed physicists.

The CMS computing system itself relies on geographically distributed resources, interconnected via high throughput networks and controlled by means of Grid services and toolkits, whose building blocks are provided by the Worldwide LHC Computing Grid (WLCG, [1]). CMS builds application layers able to interface with several different Grid flavors (LCG-2, Grid-3, EGEE, Nor-

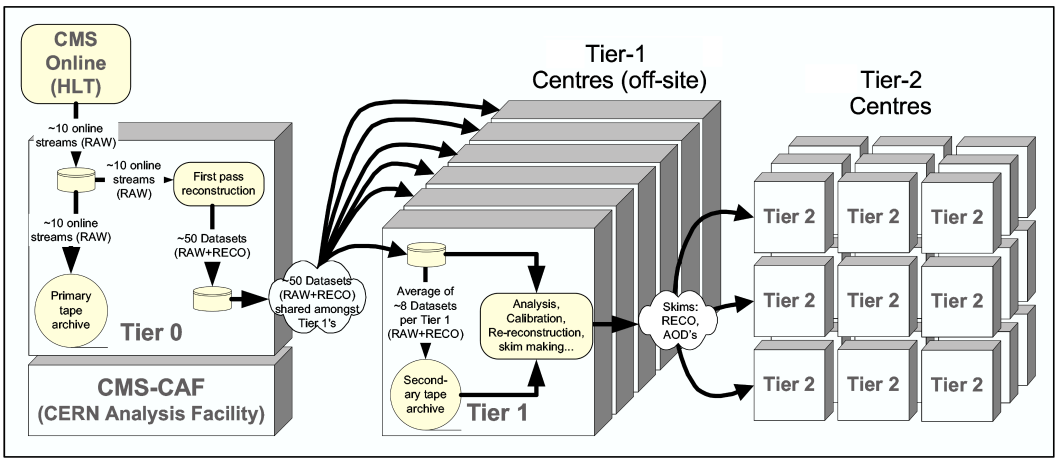

Figure 1: CMS computing model and its tiers hierarchy

duGrid, OSG). A WLCG-enabled hierarchy of computing tiers is depicted in the CMS computing model [2], and their role, required functionality and responsibilities are specified (see Figure 1).

CERN constitutes the so-called Tier-0 center: here data from the detector will be collected, the first processing and storage of the data will take place and raw/reconstructed data will be transfered to Tier-1 centers. Beside the Tier-0 center, CERN will also host the CMS Analysis Facility (CAF), which will have access to the full raw data and will be focused on latency-critical detector, trigger and calibration activities. It will also provide some CMS central services like the storage of conditions data and calibrations.

There are then two level of tiers for quite different purposes: organized mass data processing and custodial storage is performed at about 7 Tier-1 located at large regional computing centres, while a more large number of Tier-2 sites are dedicated to computing.

For what concerns custodial storage, Tier-1 centers receive simulated data produced within the Tier-2 centers, and will receive reconstructed data together with the corresponding raw data from the Tier-0. Regarding organized mass data processing activities, Tier-1 centers will be in charge of calibration, re-processing, data skimming and other organized intensive analysis tasks.

The Tier-2 centres are essentially devoted to the production of simulated data and to the user distributed analysis of data imported from Tier-1 centers. In this sense, the Tier-2 activities are much more "chaotic" with respect to the higher tiers: analysis are not centrally planned and resource utilization decisions are closer to end users which can leverage wider set of them. So the claim for high flexibility and robustness of the the workflow management leads to an extremely compound infrastructure which inevitably entails a large effort in phase of testing and deploying.

2. Workload Management System for Analysis

The Workload and Data Management Systems have been designed in order to make use of the existing Grid Services as much as possible, building on top of them CMS-specific services.

In particular, the CMS Workload Management System (WMS) relies on the Grid WMS provided by the WLCG project for job submission and scheduling onto resources according to the CMS Virtual Organization (VO) policy and priorities. Using the Grid Information System, it knows the available resources and their usage. It performs matchmaking to determine the best site to run the job and submits it to the Computing Element

(CE) of the selected site which in turn schedules it in the local batch system. The Worker Node (WN) machines where jobs run have POSIX-IO-like access to the data stored in the local Storage Element (SE).

On top of the Grid WMS, CMS has built the CMS Remote Analysis Builder (CRAB, [3–5]), an advanced client-server architecture for specific CMS-software (CMSSW) analysis jobs workflow management. It is based on independent components communicating through an asynchronous and persistent message service, which can provide for the strong requirement of extreme flexibility. Such a server is placed between the user and the Grid to perform a set of actions in user behalf through a delegation service which handles users proxy certificates. The main goal of such an intermediate server is to automate as much as possible the whole distributed analysis workflow and to improve the scalability of the system in order to fullfill the target requirement rate of more than 100 thousands jobs handled per days when LHC will be full operational.

Anyway, the client-server implementation is transparent to the end users. From this point of view CRAB simply looks like a dedicated front end for specific CMSSW analysis jobs. It enables the user to process datasets and Monte Carlo (MC) samples taking care of CMSSW specific features and requirements, provides the user with a simple and easy to use interface, hides to the user the direct interaction with the Grid and reduces the user load by automating most of the action and looking after error handling and resubmissions.

CRAB's own functionalities and its integrated interaction with the underlying Grid environment needs a dedicated test and deployment activity.

3. Test Experiences

The CMS experiment is getting ready to the real LHC data handling by building and testing its Computing Model through daily experience on production-quality operations as well as in challenges of increasing complexity. The capability to simultaneously address both these complex tasks relies on the quality of the developed tools and related know-how, and on their capability to manage switches between testbed-like and production-like infrastructures.

Such intermediate activities between developing and production are crucial due to the large number of different services running on different layers using different technologies within multiple operational scenarios that will operate when LHC will work at full. Past experience [6,7] has shown that such training activity practice is one of the biggest challenges of such a system.

Main issues concern both functionality tests as well as the scalability of the whole infrastructure to maintain the needed performance and robustness both up to the expected full job flows rates and under realistic usage conditions.

Work flow tests was performed on both on Grid level and on more CMS specific CRAB level.

3.1. Grid Workflow

WMS testing was aimed at probing the load limits of the service both from the hardware and from the software point of view. A quasi automated test-suite was set up to steadily submit jobs at adjustable rate. Some very controlled instances of WMS have been used and was continuously tested, patched and re-deployed, with a tight feedback between testers and developers.

The tests involved the submission of large numbers of jobs to the WLCG production infrastruc-

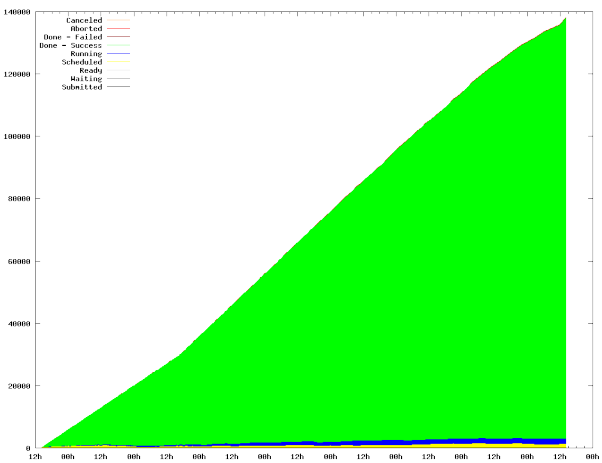

Figure 2: The Grid WMS is capable to sustain a rate of about 20kJ/d for several days

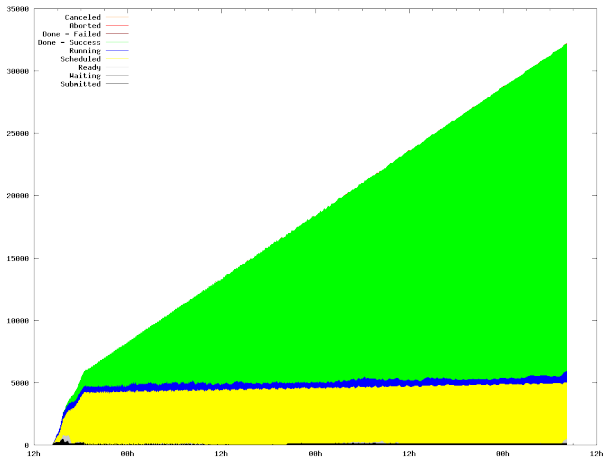

Figure 3: 5-days no-stop 10kJ/d submission on Grid CE, with always 5k active jobs

ture, both using simple "hello world" scripts and real experiment applications. Problems encountered were reported to the developers, who provided bug fixes, in an iterative process. Acceptance criteria were defined to assess the compliance of the WMS with the requirements from the CMS and ATLAS experiments and the WLCG operations: uninterrupted submission of at least 10^4 jobs per day for period of at least five days; no service restart required during this period and no degradation in performance at the end of this period, with a number of stale jobs less than 1% of the total at the end of the test.

A successful test on gLite 3.1 middleware fully met the acceptance criteria. 115'000 jobs was submitted along 7 days (16'000 job/day) with 320 (0.3%) jobs aborted due to the WMS problems, negligible delay between job submission and arrival on the CE. Further tests proved that the gLite WMS is able to sustain for a week an higher rate of 20'000 job/day without degradation in performances and no stale jobs (see Figure 2).

Testing the performance and reliability of CE was done submitting at well specified rate (10k jobs per day) in order to always keep at least 5k jobs active in the CE, according to the criteria defined for the CE acceptance tests. Figure 3 shows the results of a first 5-days no-stop submission test to verify the acceptance criteria. Something like 60k jobs was submitted; only 119 jobs aborted ($< 0.2\%$) but not due to a CE error; no performance degradation observed and the CE service never restarted.

3.2. CRAB Workflow

Testing the CRAB infrastructure is a more structured task, since the goal is now to probe not only the sustainability of a high job submission rate, but also the reliability of the whole system of services and functionalities.

For this purpose the so called JobRobot was developed: a whole agent and dropbox based automated expert system architecture which enables to simulate a massive and complete user analysis activity, from creation to submission, monitoring and up to output retrieval. The goal is to understand and solve potentially race condition criticalities that only a realistic, "chaotic" usage can lead to show up.

First very exploratory rounds of test was already performed, using an instance of a preliminary version of the CRAB server attached to a dedicated Grid WMS machine. Passing through the CRAB server, the JobRobot was keeping spreading CMSSW collections of jobs over about 30 different sites, using realistic requirements, with a growing submission rate and for a growing number of days. An initial submission rate of 10k jobs per day was quite easily sustained for 4 days, showing no overload both on the Grid

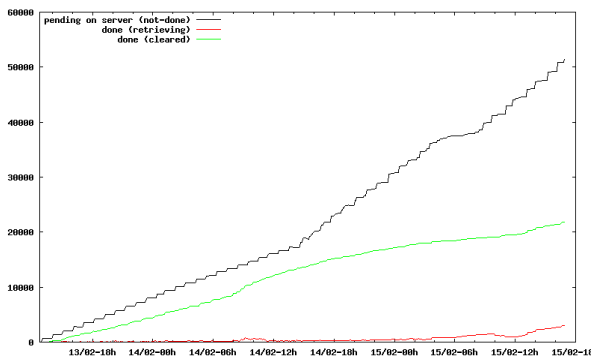

Figure 4: Starting a stress test for the CRAB server workflow at a rate of 18kJ/d

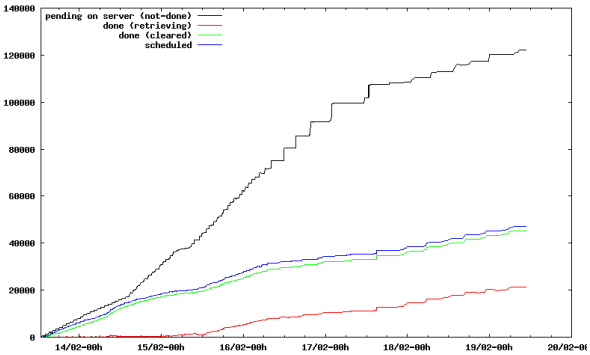

Figure 5: The bottleneck of having only one Grid WMS

WMS side as well as on the CRAB server side. In a more stress-oriented test a rate of 18k jobs per day was maintained for the first 24 hours and then raised to 30k jobs per day for the succeeding 24 hours. As Figure 4 shows, within the lower rate jobs complete their workflow (green line) at the same rate they are submitted (black line).

The high submission rate was overstretched for even more time (see Figure 5), but the single Grid WMS instance was not able to efficiently handle the overall job flow generated in this way: job was dispatched more and more slowly to sites, piling up in WMS queues and bringing a to a rapid degradation of its performace, and as a consequence of the performance of the whole system (e.g. the lowering of the black curve in Figure 5). The expected message was that to further increase the scale requires additional dedicated WMSs.

On the contrary the load on the CRAB server machine was very reasonable, proving the a single server can still fairly well handle such a high submission rate. Moreover such first tests was already capable of giving valuable feedback concerning improvements in some CRAB server components and precious feedback was provided to developers about that. As a matter of fact, it also shows that the server could represent a further testing instrument for the underlying Grid WMS services, allowing a fine tuning of its configuration parameters in a much more CMS specific use cases tailored way.

Already planned next steps involve the set up of the next forthcoming major release of the CRAB server: it includes important development upgrades (a client-server reengineering and a refactory of the Grid interaction framework) and needs an early test of the overall functionalities. A new scale test with the server pointing at more than one single Grid WMS is then scheduled.

4. Conclusion

The Grid infrastructure is already really working at production level and actually the every day CMS activity can not do without it. So far testing and integration activities represented a central part of the work needed to bring the workload management infrastructure to to a quality level which let users take full advantage of it in a production context.

Anyway there are still challenges, to be ready for when the LHC will be fully operative. This testing process is still ongoing and the improvements achieved during these months have already had a big impact on the amount of effort required to run the work flow services in production activities.

REFERENCES

1. LHC Computing Grid project: http://lcg.web.cern.ch/LCG/
2. CMS Computing Technical Design Report, CERN-LHCC-2005-023, June 2005
3. D. Spiga, S. Lacaprara, W. Bacchi, M. Cinquilli, G.Codispoti, M. corvo, A. Dorigo, F. Fanzago, F. Farina, O. Gutsche, C. Kavka, M. Merlo, L. Servoli, A. Fanfani. (2007). *CRAB: the CMS distributed analysis tool development and design*, NUCLEAR PHYSICS B-PROCEEDINGS SUPPLEMENTS. Hadron Collider Physics Symposium 2007. La Biodola, Isola d'Elba, Italy. 20-26 Maggio 2007. vol. 177-178C, pp. 267 - 268.
4. A. Fanfani, D. Spiga, S. Lacaprara, W. Bacchi, M.Cinquilli, G. Codispoti, M. Corvo, A. Dorigo, F. Fanzago, A. Fanfani, F. Farina, M. Merlo, O. Gutsche, L. Servoli,

C. Kavka(2007). *The CMS Remote Analysis Builder (CRAB)*. LECTURE NOTES IN COMPUTER SCIENCE. High Performance Computing -HiPC 2007 14th International Conference. Goa, India. 18-21 Dicembre 2007. vol. 4873, pp. 580 - 586.

5. F. Farina , S. Lacaprara, W. Bacchi, M. Cinquilli, G. Codispoti, M. Corvo, A. Dorigo, A.Fanfani, F. Fanzago, O. Gutsche, C. Kavka, M. Merlo, L. Servoli, D.Spiga. (2007). *Status and evolution of CRAB*. POS PROCEEDINGS OF SCIENCE. XI International Workshopon Advanced Computing and Analysis Techniques in Physics Research (ACAT07). Amsterdam. 23-27 April 2007. vol. ACAT20, pp. ACAT020.
6. Andra Sciabà, S. Campana, A. Di Girolamo, E. Lanciotti, N. Magini, P. M. Lorenzo, V. Miccio, R. Santinelli, *Testing and integrating the WLCG/EGEE middleware in the LHC computing*, International Conference on Computing in High Energy and Nuclear Physics (CHEP07), Victoria BC, Canada, 2-7 September 2007
7. V. Miccio, S. Campana, A. Sciabà, *Experience in testing the gLite workload management system and the CREAM computing element*, EGEE'07 International Conference. 1-5 October 2007