CMS Monte Carlo production operations in a distributed computing environment

A. Mohapatra[a], C. Lazaridis[a], J.M. Hernández[b], J. Caballero[b], C. Hof[c], S. Kalinin[c], A. Flossdorf[d], M. Abbrescia[e], N. De Filippis[e], G. Donvito[e], G. Maggi[e], S. My[e], A. Pompili[e*], S. Sarkar[f], J. Maes[g], P. Van Mulders[g], I. Villella[g], S. De Weirdt[g], G. Hammad[h] S. Wakefield[i], W. Guan[i], J.A.S. Lajas[i], P. Kreuzer[j], A. Khomich[j], P. Elmer[k], D. Evans[l], A. Fanfani[m], W. Bacchi[m], G. Codispoti[m], F. Van Lingen[n], C. Kavka[o], and G. Eulisse[p]

[a]University of Wisconsin, Madison, WN, USA

[b]CIEMAT, Madrid, Spain

[c]RWTH, III Physikalisches Institut, Aachen, Germany

[d]DESY, Hamburg, Germany

[e]Dipartimento Interateneo di Fisica di Bari and INFN Sezione di Bari, Bari, Italy

[f]INFN Sezione di Pisa, Pisa, Italy

[g]Vrije Universiteit Brussel, Brussel, Belgium

[h]Université Libre de Bruxelles, Bruxelles, Belgium

[i]Imperial College, London, UK

[j]CERN, Geneva, Switzerland

[k]Princeton University, Princeton, NJ, USA

[l]FNAL, Batavia, IL, USA

[m]Università degli Studi di Bologna and INFN Sezione di Bologna, Bologna, Italy

[n]California Institute of Technology, Pasadena, CA, USA

[o]INFN Sezione di Trieste, Trieste, Italy

[p]Northeastern University, Boston, MA, USA

Monte Carlo production for the CMS experiment is carried out in a distributed computing environment; the goal of producing 30M simulated events per month in the first half of 2007 has been reached. A brief overview of the production operations and statistics is presented.

1. INTRODUCTION

The CMS experiment is expected to start collecting data in 2008. Monte Carlo (MC) production is crucial for delivering large samples of fully simulated and reconstructed events which are required for detector performance studies, software validation and physics analysis. The production effort is carried out by using two dif-

*Corresponding author; Email: pompili@ba.infn.it

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

ferent Grid infrastructures: the American OSG (OpenScienceGrid) and the, mainly European, LCG/EGEE (LHC Computing Grid/Enabling Grids for E-sciencE). Up to 40 Tier1/Tier2 sites are shared among 6 production teams.

2. PRODUCTION SOFTWARE TOOL

The CMS ProdAgent (PA) is a multi-daemon workflow management tool [1] used to execute and manage the CMS MC production. This includes steps such as job creation, (re)submission, tracking, merging and publication of ouput files on the local/global-scope data bookkeeping, location and transfer system databases. PA has been developed with the goal of: 1) making the production as automatic as possible, adopting a tunable 'if fail try again' approach to failures; 2) handling many thousands jobs in parallel (times the number of independent istances run by each team).

A local teams typical setup includes a couple of PA istances running on a couple of User Interfaces (UI); on each UI a separate MySQL database and other client tools (integrated in PA) run.

To access the LCG resources, each PA istance submits jobs to the sites via a pool of EDG Resource Brokers (RB), working in a round robin configuration at CERN. Depending on the job configuration, the Grid site is assigned by the operator or its choice is left to the RB on the basis of a set of requirements (software version, minimum CPU time, site restriction). Direct Condor-G submission is used in OSG instead.

3. PRODUCTION OPERATIONS

Since the beginning of 2007 about $140M$ simulated events have been produced that corresponding to an average rate of about $30M$ events per month, rate value that also represents the goal for CMS computing in the first half of 2007. During recent spring production the rate even increased to $60M/45$ days. The whole amount of events is spread over a huge variety of different datasets (beyond 1100), produced with several software releases, conditions (with and without pile-up) and steps (generation and simulation first, digitization and reconstruction as a second step, and merging steps to optimize output file size for data archiving and transfer).

In 2007 the rate of MC production has been increasing thanks to the ongoing optimization of a variety of software tools (including PA), to the increase of the CMS production teams and to the collaborative effort of the personnel of the about 40 sites involved.

Overall job efficiency roughly ranges from 50% to 80% depending on the site. Grid middleware efficiency is in average about 90%. The somewhat low efficiency range shows that the production automation, a better handling of the inherent Grid unreliability and a good responsiveness of the sites administrators are crucial.

An efficient MC production, at the required production rate, is still manpower intensive and needs a lot of expertise in CMS and Grid computing; some extra manual operations or debugging duties by the teams are often required. On the other hand the wide operational feedback has been useful for 1) the production tools development towards the required level of automation, flexibility and performance, and 2) the debugging and troubleshooting of the Grid sites configuration in a common effort with the administrators.

4. CONCLUSIONS

CMS MC production system making use of distributed computing resources is fully operational. Operations teams have successfully produced $140M$ of events within five months. Thanks to the feedback provided by teams the system has rapidly been progressing towards the final design which includes dedicated PA components for increasing the efficiency. The production rate already achieved make us confident to fulfill the forthcoming CMS computing challenge and its target rate of $50M$ events per month.

REFERENCES

1. D. Evans et al. *The CMS Monte Carlo Production System: Development and Design.* Presented at the Hadron Collider Physics Symposium HCP 2007, Italy.