The CMS Monte Carlo Production System: Development and Design

D. Evans [a] *, A. Fanfani [b], C. Kavka [c], F. van Lingen [d], G. Eulisse [e], W. Bacchi [b], G. Codispoti [b], D. Mason [a], N. De Filippis [f], J. M. Hernández [g] and P. Elmer [h]

[a]Fermi National Accelerator Laboratory, Batavia, IL, USA

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

[c]INFN Sezione di Trieste, Trieste, Italy

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

[e]Northeastern University, Boston, MA, USA

[f]INFN Sezione di Bari, Bari, Italy

[g]Centro de Investigaciones Energeticas, Medioambientales y Tecnológicas, Madrid, España

[h]Princeton University, Princeton, NJ, USA

The CMS production system has undergone a major architectural upgrade from its predecessor, with the goal of reducing the operational manpower needed and preparing for the large scale production required by the CMS physics plan. The new production system is a tiered architecture that facilitates robust and distributed production request processing and takes advantage of the multiple Grid and farm resources available to the CMS experiment.

1. Introduction

The CMS experiment at CERN relies on a distributed computing model for the Monte Carlo production. This paper describes the main architectural components of the new production system.

2. The MC production system architecture

The architecture of the new Monte Carlo (MC) production system [1] consists of three components (figure 1): the Request system (*ProdRequest*), acts as a frontend application for (user) production request submissions into the production system; the Production Manager (*ProdManager*), manages these user requests, performing accounting and allocating work to a collection of Production Agents (*ProdAgent*s). The agents ask for work when resources they manage are available and manage submissions, possible errors and resubmissions while performing local cataloguing operations.

*Corresponding author. Email address: evansde@fnal.gov

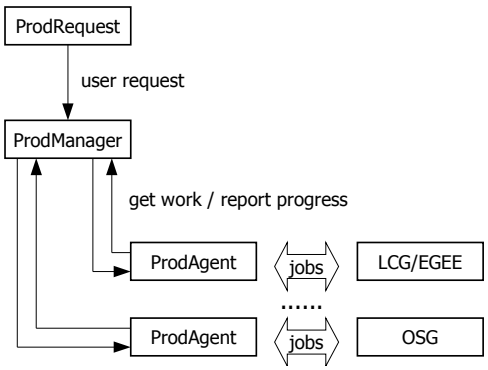

Figure 1. The structure of the MC production system.

The ProdAgents [4] themselves are defined internally in terms of autonomous components that communicate via an asynchronous and persistent message system. Components subscribe to messages of interest, get messages addressed to them waiting only if there are no messages available, and publish their messages without waiting for data to be transferred. Delayed and queued message functionality enables the ProdAgents to adequately deal with third party component interaction, like for example the CMS cataloguing system and the file transfer systems [1].

At the core of the system there is a database that registers the persistent status information like messages, job states and local dedicated data management services (DBS).

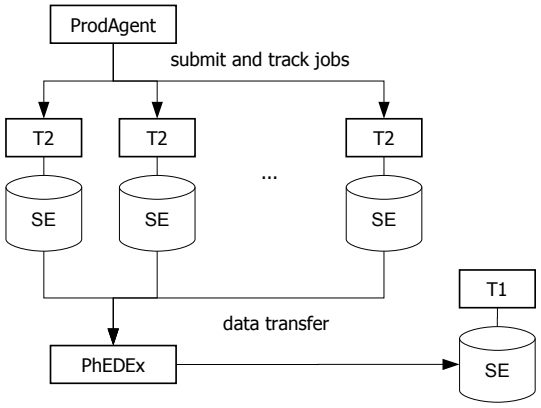

Figure 2. The job workflow model and the file transfer system.

The Monte Carlo production job workflow is shown in figure 2. The ProdAgents create processing jobs, which are submitted to the Tier2 [1] centers by using the Grid. These jobs produce output files that are stored in local storage elements and registered in the local DBS catalogs. When enough data is available in the sites, merge jobs are submitted, merging previously produced data and removing the unmerged files. The ProdAgents take care of error handling and resubmission of both production and merge jobs, triggering automatically the data transfer to the target storage system in the Tier1 centers, performing registration in global catalogs.

Various complementary monitoring tools provide end-to-end monitoring of the system to track down potential problems. A local monitoring tool provides information about the status of all components (access to log files, etc), while a global monitoring tool provides summaries of the status of the productions performed by all ProdAgents that are running in the CMS MC production system.

The framework defined by this new production system is general, and since it is structured as a set of independent components, other CMS related projects like CRABServer [3] and Tier0Agent [4] have based their structure on this new framework.

3. Conclusions

The new MC production system for CMS aims at providing a high level of automation with support for multiple Grids. The system is designed to scale to multiple request, manager and agent components thereby minimizing single points of failure. Its autonomous component based design allows many developers independently contribute to the development. The new production system has been successfully used to generate many hundreds of millions of events [2], making heavy use of LCG and OSG resources.

REFERENCES

1. *CMS, The Computing Project. Technical Design Report*. Available from the CMS web site at http://cmsdoc.cern.ch/cms/cpt/tdr
2. A. Pompili et al. *CMS Monte Carlo Production Operations in a Distributed Computing Environment*. Presented at the Hadron Collider Physics Symposium HCP 2007, Italy.
3. D. Spiga et al. *CRAB, CMS Remote Analysis Builder*. Presented at the Hadron Collider Physics Symposium HCP 2007, Italy.
4. *The ProductionAgent twiki page*. Available from CERN twiki web site.