Available on CMS information server CMS NOTE 2004/034

The Compact Muon Solenoid Experiment

CMS Note

Mailing address: CMS CERN, CH-1211 GENEVA 23, Switzerland

September 1, 2004

CMS Results of Grid-related activities using the early deployed LCG Implementations

T. Coviello[a], M. D'Amato[a], N. De Filippis[a], G. Donvito[a], G. Maggi[a], A. Pierro[a]

Dipartimento di Fisica dell'Universita' di Bari e INFN, I-70126 Bari, Italy

D. Bonacorsi[ab], P. Capiluppi[a], A. Fanfani[a], C. Grandi[a]

Dipartimento di Fisica dell'Universita' di Bologna e INFN, I-40127 Bologna, Italy

O. Maroney[a]

Bristol University, H.H. Wills Physics Laboratory, Bristol BS8 1TL, United Kingdom

H. Nebrensky[a]

Brunel University, Department of Electronic and Computer Engineering, Uxbridge UB8 3PH, United Kingdom

F. Donno[bc], W. Jank[a], A. Sciaba'[bd], N. Sinanis[a]

CERN, European Organization for Nuclear Research, CH-1211 Geneva 23, Switzerland

D. Colling[a], H. Tallini[a], B. MacEvoy[a]

Imperial College, University of London, London SW7 2BZ, United Kingdom

S. Wang[a], J. Kaiser[a]

Iowa State University, Department of Physics and Astronomy, Iowa City, IA 52242-1479, USA

A. Osman[a]

National Centre for Physics, Quaid-I-Azam University, Islamabad, Pakistan

C. Charlot[a], I. Semenjouk[a]

Laboratoire Leprince-Ringuet (LLR), Ecole Polytechnique, F-91128 Palaiseau Cedex, France

[a] CMS Experiment

[b] Also INFN LCG Project

[c] On leave from INFN Pisa, I-56127 Pisa, Italy

[d] Also at CNAF national Centre of INFN, I-40127 Bologna, Italy

M. Biasotto[a], S. Fantinel[a][e]

Laboratori Nazionali di Legnaro dell' INFN (LNL), I-35020 Legnaro (Pd), Italy

M. Corvo[e], F. Fanzago[e], M. Mazzucato[a], M. Verlato[a]

Dipartimento di Fisica dell'Universita' di Padova e INFN, I-35131 Padova, Italy

A. Go[a], K. Chia Ming[a]

National Central University, Physics Department, Chung-Li, Taiwan

S. Andreozzi[e], A. Cavalli[f], V. Ciaschini[e], A. Ghiselli[e], A. Italiano[f], F. Spataro[e], C. Vistoli[e]

CNAF national Centre of INFN, I-40127 Bologna, Italy

G. Tortone[e]

INFN Sezione di Napoli, I-80126 Napoli, Italy

Abstract

The CMS Experiment is defining its Computing Model and is experimenting and testing the new distributed features offered by many Grid Projects. This report describes use by CMS of the early-deployed systems of LCG (LCG-0 and LCG-1). Most of the used features here discussed came from the EU implemented middleware, even if some of the tested capabilities were in common with the US developed middleware.

This report describes the simulation of about 2 million of CMS detector events, which were generated as part of the official CMS Data Challenge 04 (Pre-Challenge-Production). The simulations were done on a CMS-dedicated testbed (CMS-LCG-0), where an ad-hoc modified version of the LCG-0 middleware was deployed and where the CMS Experiment had a complete control, and on the official early LCG delivered system (with the LCG-1 version).

Modifications to the CMS simulation tools for events production where studied and achieved, together with necessary adaptations of the middleware services. Bilateral feedback (between CMS and LCG middleware) played an important role in making progress (including bugs corrections).

Fractions of successful processing of simulation jobs ranged from 70% to 90%. Most of the failure reasons were identified, with RLS instability being the greatest cause of failure.

Evaluation of the LCG-1 middleware is also presented and discussed. The progress of the new functionalities introduced and the better-distributed organization of services were tested and eventually stressed. While the overall efficiency decreased in the early implementation of the system in respect of the LCG-0 testbed, a consistent success rate of 50-60% was achieved. One of the major difficulties for the simulation of CMS events on the LCG-1 system was identified to be the consistent configuration of the distributed sites.

---

[e] EU DataTAG Project

[f] INFN-CNAF LHC "Tier1" Project

Table of Content

1 Introduction

The Compact Muon Solenoid experiment (CMS) [1] is one of the four particle physics experiments that will collect data at the Large Hadron Collider (LHC) [2] being built at CERN (Geneva, Switzerland) [3]. The preparation and building of a Computing System able to treat the data that will be collected by the experiment has to go through steps of increasing complexity before its final implementation. The Data Challenge during 2004 (CMS DC04) was one of the planned steps towards producing the CMS distributed System and required a number of developments in the software, the middleware, the organization and the distribution of tasks among the CMS Regional Centres.

Here we present the work done in adapting and testing the CMS software for the EU Grid environment and in particular the LCG (LHC Computing Grid) Project [4] in the context of Pre-Challenge-Production for DC04. Approximately 2 million events were generated using CMS tools integrated with the early LCG system (LCG-0). An evaluation of the LCG-1 middleware was also performed, in order to continue the migration and running of CMS Data Challenge on the new environment (waiting and preparing for the LCG-2 system deployment).

This work demonstrated the importance of close collaboration between many of the actors of the current LHC Computing effort, with focused goals (such as an Experiment Data Challenge), in particular the EU-EDG Project [5], the EU-EDT Project [6], the LCG Project and the CMS Experiment developers.

The CMS DC04 originally consisted of the fake running of the CMS off-line data treatment for a period of one month at LHC "low Luminosity". To achieve the challenge a ~50 million events had to be simulated and digitized (raw data) in advance: the so-called Pre-Challenge Production (PCP), starting from July 2003 and foreseen to last about six months. The Data Challenge itself ran at the beginning of 2004 (hence DC04).

The Grid-related activities reported here are parts of the PCP official CMS activity and had two main phases:

- A preparation phase and testbed set-up, necessary to adapt the CMS production tools to the Grid environment and also to better understand (and possibly modify) the modularity of the tools themselves.
- A real production phase during the PCP time, aiming to produce as many events as possible on the deployed distributed grid system.

This report summarizes the use of many contributions of middleware implementation/adaptation, tools development and services deployment/maintenance. Actors of these activities were the EU-EDG Project, the EU-EDT Project, the LCG Project and the CMS Experiment, with different but complementary roles.

2 Integration of CMS production tools with LCG middleware

The CMS production tools for job creation, job submission and job tracking are collected under the CMS OCTOPUS project [7].

Figure 1 shows the design for interfacing the CMS Production system to LCG Middleware, which were then integrated into the LCG releases.

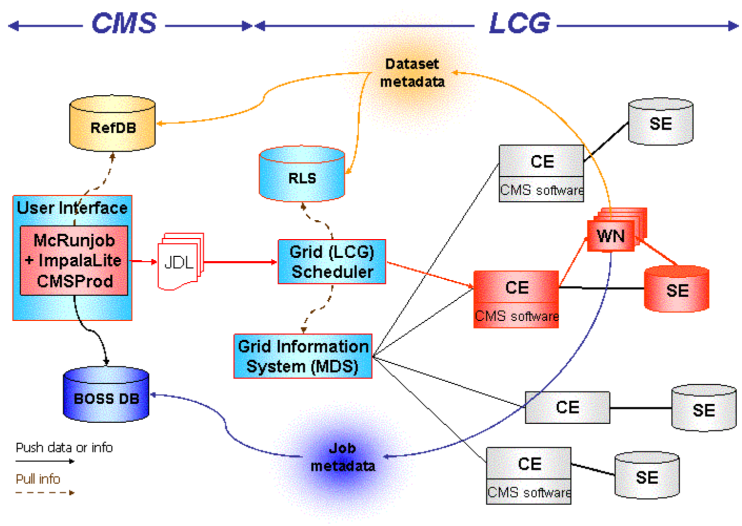

Figure 1: Integration of CMS production tools with LCG middleware.

The main components of CMS production are:

RefDB [8]**,** a central SQL database, located at CERN, which allows the gathering of all relevant information to be given as input to the production component (McRunJob or CMSProd) that takes care of starting a specific production "assignment". The RefDB also keeps the summary information about the status of each assignment.

McRunJob (Monte Carlo Run Job) [9], a Python tool for job preparation that provides a metadata based approach for specifying workflow patterns and that is modular. It has plug-in's for submitting jobs to several resource management systems.

CMSProd [10], a Python program intended to simplify the process of creation and submission of CMS production jobs. The program is able to generate scripts for different job types and submit them to several batch systems.

SRB (Storage Resource Broker) [11], a client-server middleware that provides a uniform interface for connecting to heterogeneous data resources over a network and accessing replicated data sets. SRB, in conjunction with the Metadata Catalog (MCAT), provides a way to access data sets based on their attributes rather than their names or physical locations. Every CMS Regional centre used SRB to download and upload data in order to make them available to the whole CMS production team.

McRunJob or CMSProd have to be installed on the UI, where the job preparation and job submission take place. Both McRunJob and CMSProd were also implemented to run production jobs in an LCG-like environment [12]. The set-up of the CMS environment appropriate to run the job is defined relying on the pre-installed CMS software in the Computing Element, as advertised in the Information system. The LCG data management tools were used to discover the physical location of an input file and to transfer the produced output data from the worker node to the Storage Element, registering it in the RLS. Along with the job scripts, LCG JDL files were created with all the scripts and parameters needed by the running job sent through the *input sandbox*.

The jobs were submitted through the **BOSS** (Batch Object Submission System) [13] interface to LCG, keeping the submitted jobs information via the BOSS real-time monitoring and bookkeeping.

The BOSS interface to the LCG system supports multiple configuration files, to allow an easy switch of Resource Broker in the case of one of them being unavailable. LCG related information (i.e. Logical File Names of the output files as registered into the RLS, status of the replica manager copy and registration, etc.) about the running job are monitored and collected in the BOSS MySQL database.

The summary of job tracking done by McRunjob/CMSProd using BOSS-filtering scripts was sent by the running job to update the RefDB.

3 The CMS/LCG-0 testbed

The main goals of the CMS/LCG-0 testbed were:

- demonstrate the usability of Grid implementations for large official CMS productions;

- provide early feedback about the necessary modifications of tools (and eventually of design) both to CMS and to LCG.

The main actors of this implementation were the CMS experiment and the LCG EIS service, from where most of the effort came. Also EDG people, EDT experts and local site managers gave a substantial contribution to support the CMS PCP-production on the deployed Grid system.

3.1 Middleware and services

The middleware installed in the testbed consisted of the following basic components:

- Virtual Data Toolkit (VDT) 1.1.6 [14], a set of software containing, among other things, Globus 2.0 and Condor 6.4.3, and modified to support the GLUE information schema;
- The release 1.4 of the EDG middleware, modified to support both the Globus and EDG implementations of the Replica Location System (RLS), the GLUE schema and the VOMS (see paragraph 4.2); it included the version 2.0 of the EDG Replica Manager;
- Virtual Organization Management System (VOMS) version 1.1.11 [15];
- The GLUE information schema, version 1.1, with extensions for the monitoring [16];
- Globus RLS, version 1.1 patch 6;
- EDG RLS, starting from version 2.0.4 and subsequently updated;
- GridICE monitoring system by EDT [6];
- Storage Resource Broker (SRB) version 2.1.2, with CASTOR support [17].

The services used by the testbed were:

- a VOMS server;
- workload management system services (resource broker, logging and book-keeping server, etc.);
- GridFTP servers on the storage elements;
- Local Replica Catalog servers, or LRCs (components of the RLS), both in the Globus and the EDG flavors;
- a Replica Metadata Catalog server, or RMC (component of the EDG RLS);
- a SRB server using CASTOR at CERN as storage resource.

The SRB was used to permanently store the generated Monte Carlo events and make them available to the CMS collaboration, and it was not interfaced to any other software. The modified EDG middleware was not an official implementation and delivery of the EDG project.

3.2 Testbed deployment

The middleware deployment was based on the LCG-0 software distribution, released in February 2003, with modifications to reflect the most recent features and developments in the middleware (see http://cmsdoc.cern.ch/cms/LCG/LCG-0). Every new modification triggered a new updated release of the installation distribution and related instructions. This procedure helped to achieve a complete and consistent set of scripts and how-to's, however it made sometimes difficult to trace the origin of problems and mis-configurations on different sites due to unaligned versions of the used scripts.

Set-up of testbed basic services

The Resource Broker used is "DataTAG/LCG-0" broker based on the EDG broker version 1_2_21_2 from EDG 1.4.3, modified to implement a complete interaction with GLUE Schema-based data from the Globus MDS. Similarly, the job description language (JDL) was extended to support the attribute VirtualOrganisation (to specify the name of the Virtual Organization, as required by the VOMS), and the broker was hence modified for a VO-based matchmaking. In addition, the support of the interaction with the EDG RLS was achieved, allowing it to understand values of the JDL ReplicaCatalog attribute of the form rls://hostname:port and edgrls://hostname:port.

The LRC was used to store GUID-PFN associations and the RMC to store LFN-GUID associations. The EDG

Replica Manager only supported the LDAP-based Replica Catalog and had to be modified, as well as some Globus libraries.

Two Resource Broker machines were installed in the testbed, one at CNAF and one at CERN, the latter having the full development environment and being used for building and testing new features in the middleware until the required robustness was achieved. The broker services were tested using a test-suite (http://cmsdoc.cern.ch/cms/LCG/LCG-0/test-suite/) which exploits customizable fake jobs, thus allowing job submission from a UI using a predefined RB, and also tests the copy of the data to the CloseSE and the registration to the RLS. The response of the overall system to different frequencies of jobs submission to the RB was therefore tested. Some instability in the JSS system was observed when submitting more than 160-200 concurrent jobs. An acceptable robustness and stability of the broker services were achieved, and then both RBs (CNAF and CERN) were ready to be used for official production activities.

A dedicated machine running the GridICE monitoring server was installed at CNAF and the EDT cluster monitoring and information providers system was deployed on the testbed. A *Glue* extended schema was installed on each resource. The installation of EDG WP4 information providers for the CE on top of the EDG software was done manually at each site according to instructions provided in collaboration with Glue experts.

A Globus MDS-based Information Service was deployed. The central Index Services have been configured at CNAF. The Information Service was structured according to the EDG 1.4.x style, with a MDS GRIS/GIIS hierarchy, a TOP GIIS and a BDII (Information Index), whose roles are outlined below.

The **GRIS** resides on each CE and SE, produces information about the status of the resource, and stores it in LDAP (ldif) format. The GRIS registers to a local **GIIS**. There is one GIIS per site, on the CE, to collect and cache information of the site nodes. The local GIIS registers itself on a TOP GIIS, which is configured on CNAF CE, nevertheless served the entire testbed by collecting and caching information from all involved sites. The Berkeley Db Information Index (**BDII**), installed at CNAF, was used to provide the RB with a fast and stable caching of information about the Grid resources, helping to avoid the known instability of the Globus GIIS. It is a standard OpenLDAP server, which periodically (every 10 minutes) feeds itself with the content of the TOP GIIS.

More details about the middleware components deployed in the testbed may be found in [22]

Table 1 lists the location and names of the machines providing the EDG Services.

EDG Service	**Used by**	**Site**	**Name**
Resource Broker	CNAF, Bari UI	CNAF	testbed010.cnaf.infn.it
Resource Broker	PD, POLY UI	CERN	cmslcgco03.cern.ch
Replica Location Service	Jobs on WN and UI	CERN IT	rlscms.cern.ch
DBII	CERN and CNAF RBs	CNAF	giis.cnaf.infn.it
Virtual Organization server	all users	CNAF	dell19.cnaf.infn.it
GridICE server	Users	CNAF	gridice.cnaf.infn.it

Table 1 – EDG Services location and names.

Addition of a site in the testbed

The addition of sites in the system was done gradually once they were checked to be operational. First of all, every interested site was requested to have as minimal resources availability a machine acting as a CE and at least one WN. A machine acting as SE and one as UI were installed, depending on the availability on each site. The set-up of a site required the NFS sharing of the `/home` directory of the CE to WNs and SE, and of the storage directory of the SE to CE and WNs (actually CE is needed only if it also plays a WN role). Attention must be put in the (Linux) UID and GID of users involved in the job submission, to avoid permission problems on mounted file systems. The CMS/LCG-0 middleware distribution could be downloaded from the official repository at http://cmsdoc.cern.ch/cms/LCG/LCG-0, for CE, WN, SE and UI nodes. Concerning the Grid certification, host X509 certificates had to be requested to the Certification Authority (CA) relevant for each site, and had to be installed on CEs and SEs. In addition, the certificates of Certification Authorities must reside on each grid node as well. In a few cases, it was also needed to include Certification Authorities that EDG was not yet aware of. Next step was the deployment of the MDS-based Information Service. At the beginning, a set of configuration files was distributed together with instructions about the customization for each site. Further on, these directives were directly included in the official installation kit. This increased the effort required by the CNAF and EIS team but allowed to reduce the amount of hands-on intervention requested on each site.

After the installation and configuration steps were completed on a site, the CNAF personnel verified step-by-step all the functionalities:

- a match-making test to show the site availability from the RB point of view;
- the submission of simple JDL to verify the basic functionality of the gatekeeper/scheduler;
- the copy of files from/to the SE;
- the check of consistency among the MDS content and the actual site configuration via ad-hoc LDAP queries.

The sites that were involved in the aforementioned steps were: Bari, Bologna, Bristol, CERN, CNAF, Ecole Polytechnique, Imperial College, ISLAMABAD-NCP (at CERN), Legnaro, Milano, NCU-Taiwan, Padova and University of Iowa. Some sites only participated to the aforementioned middleware deployment phase, and did not take part in official production activities over the testbed system, as outlined below; they provided very useful feedback about the deployment procedure and the reliability of the installation kit.

The CMS software was distributed as RPMs, prepared from the tar distributions used in official CMS productions. The correct CMS software installation at several sites was checked submitting to each site CMS validation jobs. When the testbed became usable for official CMS production, a brand new CMS regional center name (and relative contact persons) was defined.

The CMS/LCG-0 testbed was considered as a *virtual* Regional Centre, grouping several CMS sites, and it received production "assignments" as any other local CMS Regional Centre.

The resources used for the official CMS PCP are summarized in Table 2. During official productions, four UIs were active at: Bari, Bologna, Padova and Ecole Polytechnique.

They managed the job submission effort, each relying on a local BOSS database. During production, the number of UIs was reduced to two, allowing reducing the needed manpower. The information contained in the BOSS databases were read and processed to produce statistics plots using the boss2root [18] tool.

Site CE	**Number of CPUs**	**CPU Type**	**CPU clock (GHz)**	**Memory (MBytes)**	**CE SI2000**	**SE Disk Space**
CERN	20	P IV (Xeon)	2.8	1000	19680	1.4 TB + 700 GB (+ 700 GB SRB)
Bari	2	P III	1.0	1000	11660	270 GB
	4	P III	1.133	1000		
	12	P III	1.266	1000		
Bologna	20	P III	1.0	500	11070	900 GB
	2	P IV (Xeon)	2.4	1000		
Legnaro	22	P III	1.0	500	34000	370 GB
	28	P IV (Xeon)	2.4	1000		
Padova	18	P III	1.0	500	33300	430 GB
	14	P III	1.266	1000		
	18	P IV (Xeon)	2.4	1000		
Ecole Polytechnique	5	P III	0.8	500	1830	220 GB

Table 2 – Computing and storage resources used during the official CMS Production. Details are provided on the PC's used (CPU type, clock and memory). No hyper-threading was used on the P IV CPU's. The SpecInt2000 relative to each farm was calculated assuming a value of 846 SpecInt2000 for a 2400 MHz P IV and the relative performance of the other CPU's as measured during the production (Figure 3).

3.3 Description of the Production and results

CMS job description and workflow

The jobs used were part of an official CMS effort, the PCP, aimed as said before at the production of ~ 50 Million simulated and digitized events at the different CMS Regional Centres.

The CMS Monte Carlo production chain performed within the CMS/LCG-0 virtual Regional Centre consisted in two steps:

- **Generation of physics events (CMKIN).** The CMKIN program requires some control "cards" (directives) as input to define the physical process (*dataset* in the HEP terminology) to be simulated and a CPU time to generate one single event that depends strongly on the physical process under consideration. The output is a random access Zebra file (*ntuple*) with only a small dimension of output per event.
- **Simulation of tracking in the detector (CMSIM).** The CMSIM program requires the input file produced by a corresponding CMKIN job and some additional input info "cards". The CMSIM step is very CPU intensive, and with a non-negligible I/O requirement.

The computational characteristics of CMKIN and CMSIM steps are summarized in Table 3, where the processing time and the corresponding output size are reported, for the two datasets assigned to CMS/LCG-0. Both CMKIN and CMSIM steps for dataset "*mu03_MB*" were processed while only CMSIM production was done for dataset "*mu03_bb2mu*". The CMKIN step for dataset "*mu03_bb2mu*" was performed in another regional centre in which there weren't enough resources to complete the CMSIM step, so the output ntuples were made available in SRB and the corresponding CMSIM phase assigned to CMS/LCG-0 testbed. In order to make the ntuples of dataset "*mu03_bb2mu*" available on the Grid, they were downloaded from SRB, stored in the CERN Storage Element and registered in the CMS RLS.

In Figure 2 is shown the distribution of the "time per job" for several CPU's relative to the CMSIM production of dataset "*mu03_bb2mu*" and in Figure 3 is shown the relative performance of the different CPU's in the same production. The different CPU's have a uniform behavior over the CMSIM productions.

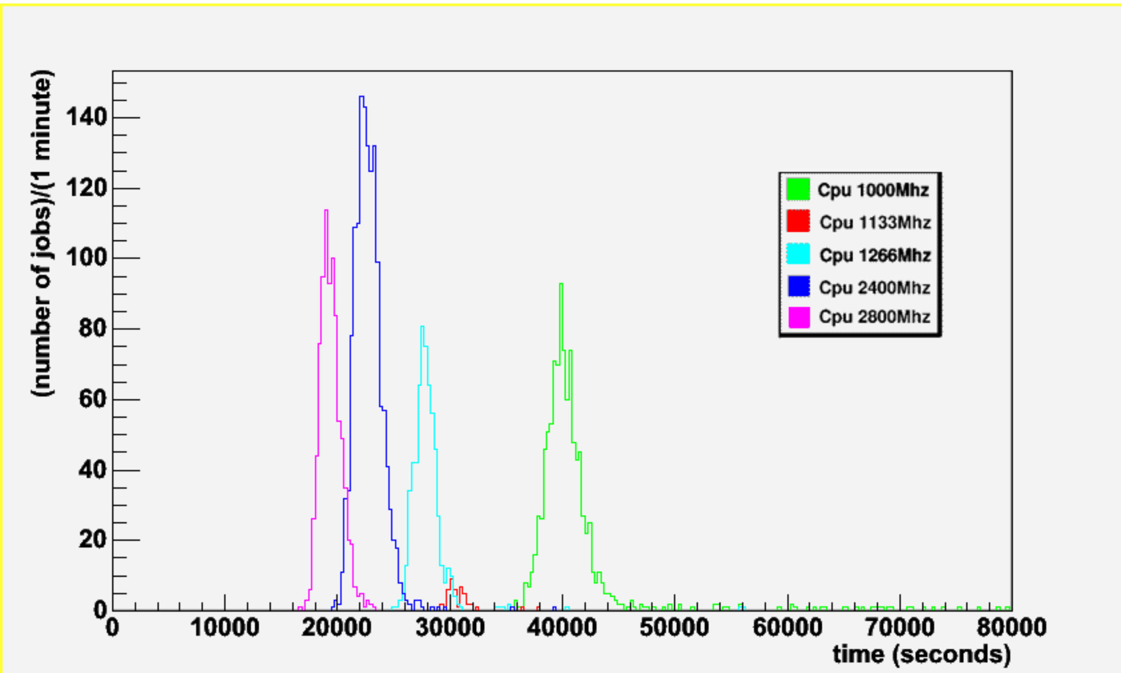

Figure 2: Distribution of the "Time per job" for different CPU types during the *mu03_bb2mu* dataset CMSIM production.

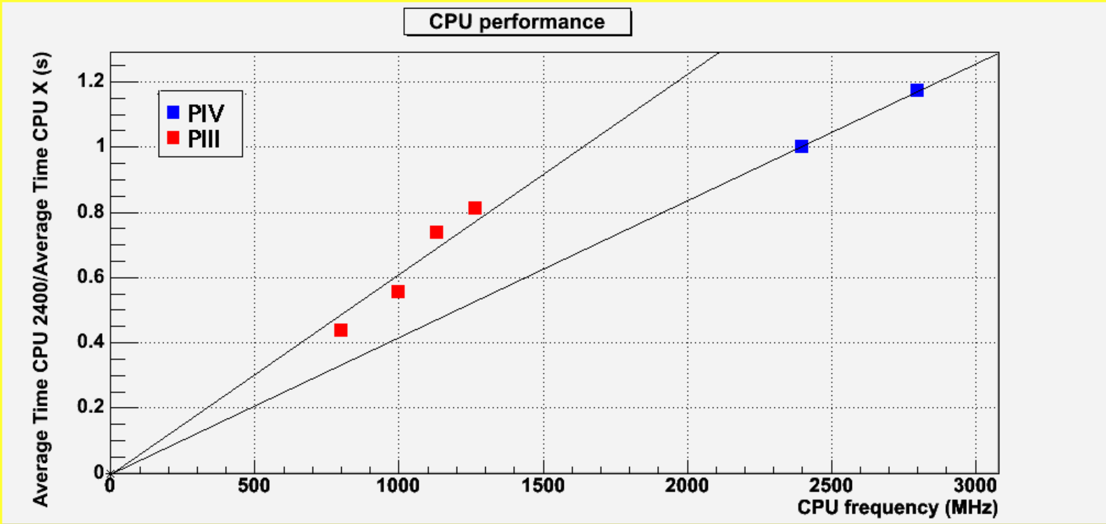

Figure 3: CPU's relative performances as a function of the clock rate during the CMSIM production of the *mu03_bb2mu* dataset. The performance of the Xeon 2,4 GHz, 1GB was assumed equal to one.

The available resources were matched against the Resource Broker, according to the job requirements to only run on the sites with CMS software installed and validated, which also provide enough "long" running queues.

The size of the CMKIN output job was small (~ 10 MB) so that the most efficient solution was to copy the needed input data locally instead of forcing the CMSIM jobs to run close to the data. The replica manager was used to discover the physical location of the *ntuple* files via its Logical Filename, and the input data transfer to the worker node (where the job was assigned to run) was done with *globus-url-copy* using the *gridftp* protocol.

The job output data was stored in the SE close to the CE where the job ran, and registered to the EDG RLS. These files were later-on replicated with the EDG Replica Manager on a SE at CERN, installed on a machine which hosted also a SRB client and server. The SRB server managed two SRB *resources*, one disk-based (coinciding with the SE host) and another one CASTOR-based.

In order to avoid filling up the disk space in the different SE's, the CMSIM output already replicated at CERN were removed from their original location. Input files to CMSIM jobs produced by other CMS regional centers were made available to the CMS/LCG-0 testbed through the SRB client at CERN and the CERN-SE.

	CMKIN				**CMSIM**			
Dataset	Time*/event	Size/event	Time*/job	Size/job	Time*/event	Size/event	Time*/job	Size/job
mu03_MB	~ 70 sec	~0.04 MB	~4.6hours	10 MB	~ 90 sec	1 MB	~6.3hours	250 MB
mu03_bb2mu	-- +	~0.04 MB	-- +	10 MB	~ 90 sec	1 MB	~6.2hours	250 MB

Table 3 – Size of data sample and CPU time per event simulation of CMKIN and CMSIM production steps (* *Xeon 2.4 GHz, 1GB*), for the two dataset that were simulated within the CMS/LCG-0 testbed. The CPU time and size per job, each job processing 250 events, is also reported. Time per event and per job on a different CPU can be obtained using the CPU relative performance shown in Figure 3.

+ *This CMKIN step was produced elsewhere and its ntuples were sent to CMS/LCG-0 to perform the CMSIM step.*

Results and phenomenology of problems

About 2 Million events of the official CMS production were produced, as reported in Table 4, corresponding to about 1.5 TB of data.

	Total Nb. of events	Size of data
CMKIN	500.000	20 GB
CMSIM	1.486.500	1.500 GB

Table 4 – Total number of successfully produced CMKIN and CMSIM events and the corresponding size of stored data

The total number of successful 'long' (CPU intensive) jobs of 250 events each was about 8000. A snapshot of the number of jobs in the system over a period of one month, according to GridICE monitoring, is shown in Figure 4. The inactive week was not due to grid instability but mainly to the set-up of the CMSIM step, in term of software distribution and modifications of CMS production tools.

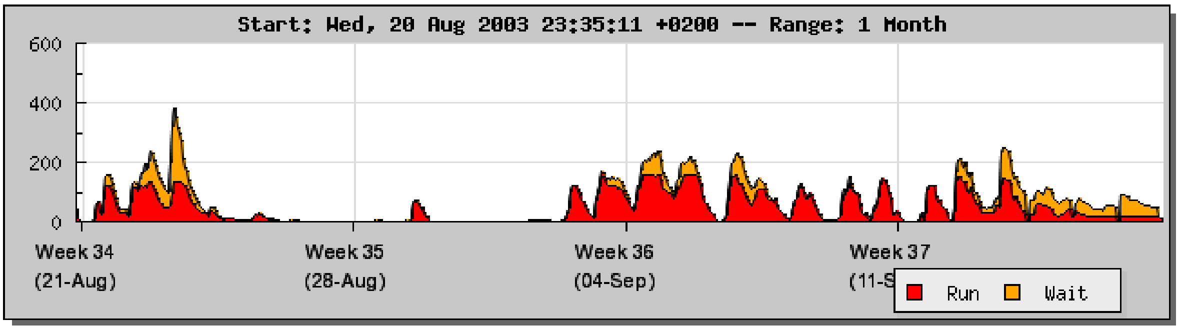

Figure 4: Number of jobs, both running and waiting, as a function of time, as shown from GridICE monthly graph. The inactive week was due to the preparation phase for CMSIM step.

A chronological summary of the integrated total amount of CMKIN events produced is shown in Figure 5. The slow activity during the first weeks of August 2003 was due to people taking care of job submission being in holidays.

In Figure 6 a chronological summary of the integrated total amount of CMSIM events produced is shown. The plateau around end of September was due both to the preparation phase for a new assignment, downloading the needed input ntuples from other sites, and to the general electric power black out in Italy which caused some services to be unavailable for few days.

The distribution of the jobs over the Computing Element is shown in Figure 7 and Figure 8, for CMKIN and CMSIM respectively. The plots show the use of the available resources, with more jobs sent to sites with more powerful resources. The difference between the number of CMKIN and CMSIM jobs executed at Legnaro is due to the fact that the resources were increased from 4 to 50 CPUs just before the start of CMSIM production.

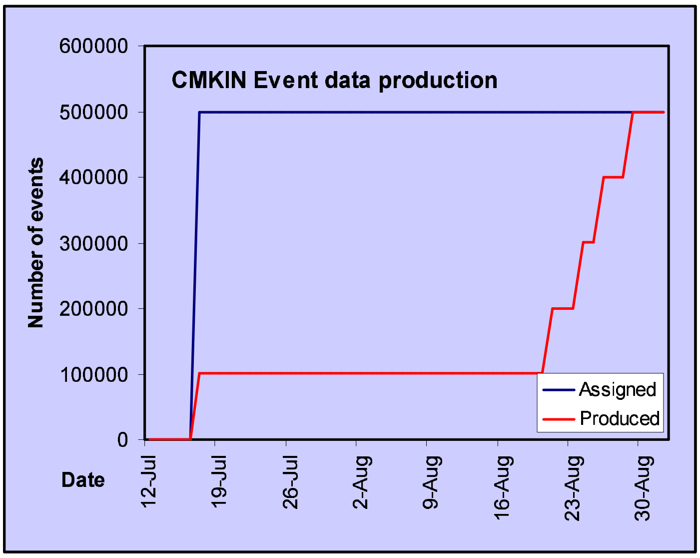

Figure 5: Chronological summary of the integrated amount of CMKIN events produced (from RefDB information, with 100k events resolution) within the CMS/LCG-0 virtual Regional Centre.

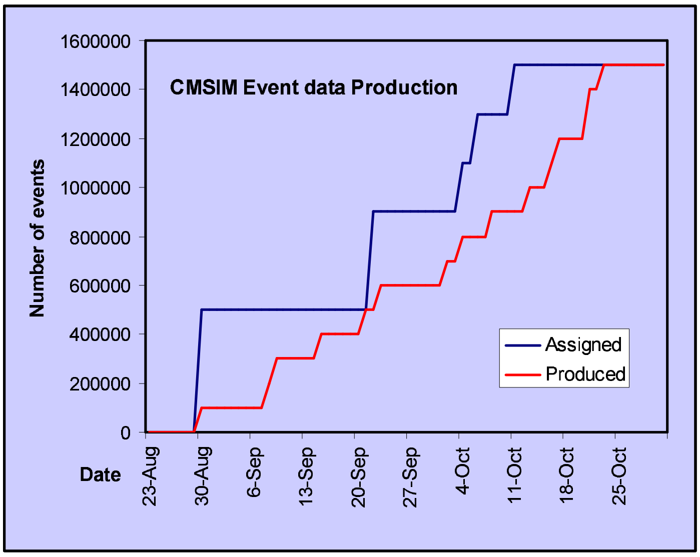

Figure 6: Chronological summary of the integrated amount of CMSIM events produced (from RefDB information, with 100k events resolution) within the CMS/LCG-0 virtual Regional Center

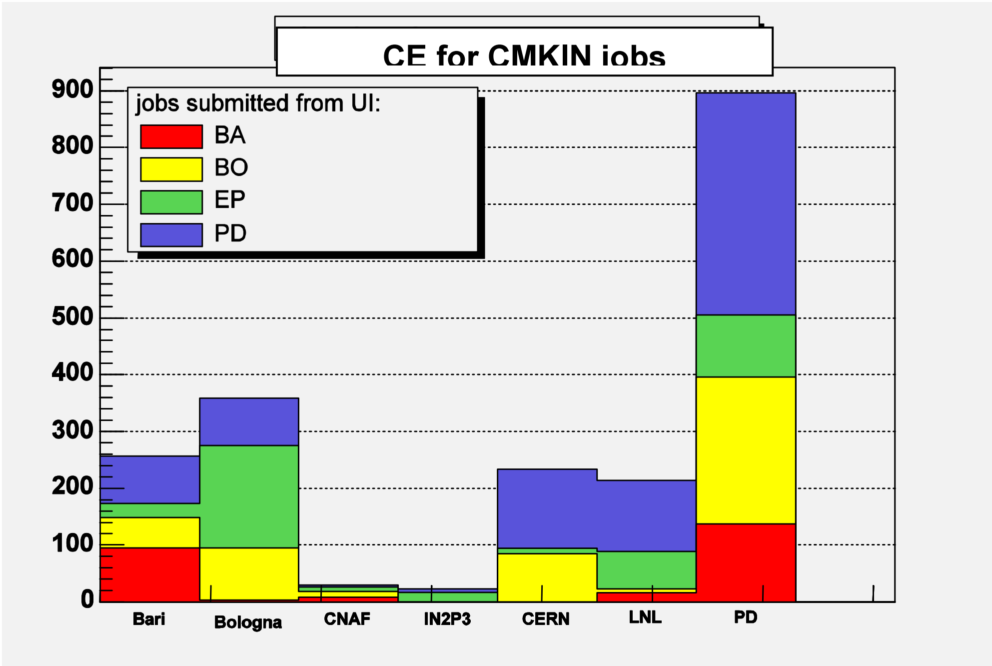

Figure 7: Number of CMKIN jobs that where executed on each CE.

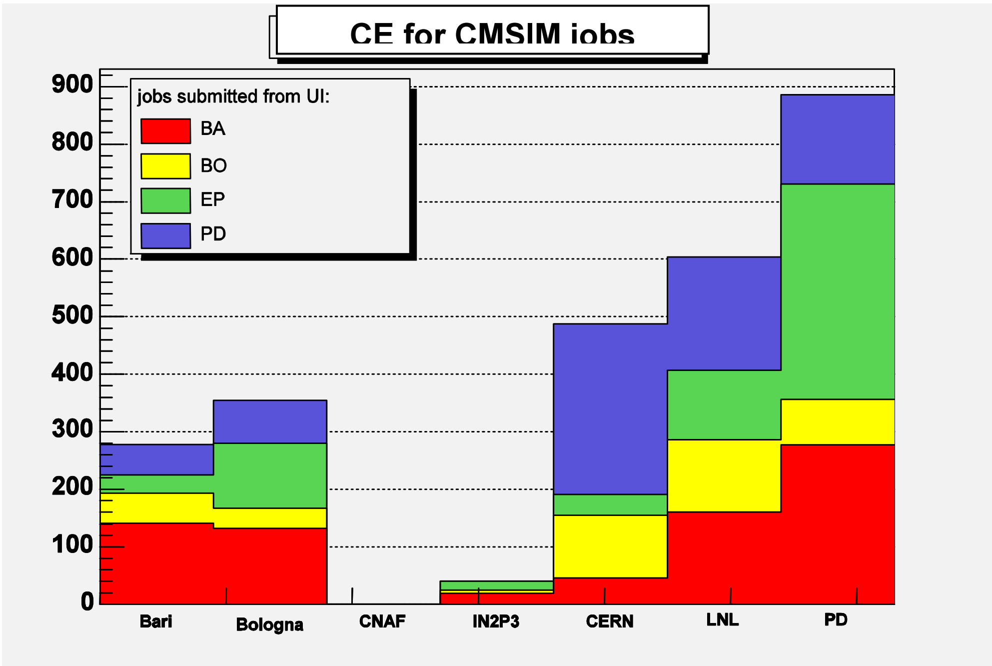

Figure 8: Number of CMSIM jobs that where executed on each CE.

A total of 500000 CMKIN events for dataset *mu03_MB (a minimum-bias sample of events needed for the muon analisys)*, a time consuming generation of un-weighted minimum bias events with muons, were produced with an efficiency at the level of 80%.

A total of 1.486.500 CMSIM events for datasets *mu03_MB* and *mu03_bb2mu (a sample of events with two muons in the final state)* were produced with an overall efficiency about 77%. The overall failure rate varied from 5-10% to 25-30% depending on the incidence of some problems, in particular those related to the RLS unavailability.

The main sources of trouble were identified as:

- The RLS unavailability, causing problems to the jobs that have just started and to those that are about to finish. At the very beginning of a job, if the RLS query fails, the physical location of the input files is not found and the job itself fails. At the end of the job the output file is registered into the RLS and the registration hangs if the RLS is not available. The latter is a very difficult situation to deal with, since the jobs stay in "Running" status forever, thus requiring manual intervention from the local system administrators. The RLS happened to be unavailable a couple of times and during these periods the job failure rates could increase up to 25-30%. Most of the problems were caused by some unexpected problems, for example disk failure and bug in the Oracle application server. In addition, some RLS intervention occurred and the experiments were usually informed about that one day in advance. However in a sustained production with the goal to saturate the available resources with "long" jobs (lasting from 7 to 15 hours) it was not always possible to empty the queues before the intervention and it was also not convenient in term of efficient usage of the testbed. Due to the fact that RLS is a single point of failure, even planned intervention can't completely prevent observing job failures during a production period.

- Some instability due to a variety of reasons at various sites (from mis-configuration, to network problems, hardware failures). These problems gave an overall inefficiency of about 5-10%. Some examples are:
    - Several jobs were failing due to "Globus down" in the Bologna CE when it was under heavy load (all CPUs busy and several jobs waiting in the queue). It turned out that the CE connections to the RB were filtered by the firewall of the Bologna's Physics department. The range of port number enabled from the CE to the RBs on the local firewall was therefore increased.
    - The home partition of Legnaro CE was getting full because all the jobs had the working directory on the home partition and with 50 running jobs writing about 250MB each, the home partition was saturated. The home size was therefore increased.
    - Problems with the LRMS (PBS at Bari and Bologna, LSF at Padova)
    - Network problems and power cut at several sites (CERN, INFN sites, LLR), in addition to a general electric power black out in Italy.
    - The CNAF RB had to be rebooted because of misbehavior caused by a disk error.
- The error "*failure while executing job wrapper*" occurring at a very low rate (few %). There could be many possible reasons behind this behavior (problems in downloading/uploading the sandbox, home directory temporary not available on WN, "qstat" transient failures for PBS, etc).

A much more improved situation was found with respect to the Stress Test [19] production carried out during last year, in term of stability and efficiency for 'long' jobs.

Some of the problems faced with where also related to the CMS production and tools:

- Problems in processing some input ntuples because they were either corrupted or their *cksum* was not consistent with the value in RefDB. These problems where reported in Savannah [20], the bug tracking system used for OCTOPUS, and they are under investigation within the CMS production team. This effect is anyway below 1%.
- Problems trying to trace failures of a single job, in particular failures in the RLS registration and in updating RefDB, and determine weather a job has to be resubmitted or not. These problems were related to lack of a robust tracking system and a resubmission mechanism within the tools for job preparation and submission.

4 Evaluation of LCG-1

The CMS/LCG-0 testbed was setup waiting for availability of the LCG-1 service; therefore as soon as the LCG-1 environment was deployed and open to the users, CMS people started to test it. The aim was again to possibly use the new service for CMS adaptation of production tools, PCP official simulation and early evaluations for Data Challenge's CMS analyses.

The porting of CMS production tools were experimented in the "Certification & Testing" LCG service, which was better suited for this kind of activities and made available earlier.

The tests of "production like" CMS jobs were launched on the LCG-1 delivered system, so as to test the first largely distributed service officially deployed by LCG.

4.1 CMS tools porting on LCG-1 environment

The LCG-1 Certification and Testing testbed was used with the goal of migrate the CMS production tools on

LCG-1 environment. The Italian Grid (*Grid-it* [21]), based on LCG-1 installation, was also used for this purpose.

The following modifications to the CMS tools were needed:

- In the tool for job preparation (McRunJob and CMSProd) the set of commands that from a running job allow staging in/out data from/to an SE and registering them into the RLS, were modified to use the new data management tools. This lead to some simplification, for example the creation of subdirectories into an SE and checking of their existence is automatically done by the *copyAndRegisterFile* command. Some modifications were required also to cope with some differences in the JDL syntax.
- BOSS interfaces to the EDG WMS through a set of scripts that were modified to use the new WMS commands to submit jobs, check their status and cancel them. The switching between Resource Brokers is obtained by specifying a configuration file with the -config-vo option.
- The filtering script to extract from standard output relevant jobs' information to be stored in the BOSS database was tuned on the data managements tools used. The printout messages of the replica manager commands available within LCG-1 were much improved with respect to the corresponding commands in EDG 1.4. In addition to the information that were extracted when using the EDG 1.4 data management tools, the time spent during the copy and registration of the output file and both the GUID and the LFN of the registered file were also stored in BOSS.

During the testing of these modifications several CMS jobs, typically lasting at maximum 30 minutes, were submitted and ran smoothly.

4.2 Tests and results of CMKIN and CMSIM jobs

The CMS software was distributed at all sites of the LCG-1 production facility during the LCG-1 middleware upgrade that was completed at the end oft October 2003. Since then CMS was allowed to start testing the submission of large number of jobs spread across several sites.

The test-jobs were of a similar kind of those used for the Pre-Challenge Production done in the CMS/LCG-0 testbed (as described in 3.3). However the submitted jobs were not part of the CMS physics production, but just for testing. The physical process generated in the CMKIN phase required a CPU time of the order of few minutes for a single job, the CMKIN job was therefore a "fast" job with small I/O (~10 MB output only). The *ntuples* produced in the CMKIN step were retrieved and processed during the CMSIM step, where the CPU intensive simulation of tracking in the detector was performed, requiring of the order of 7 hours for job (on a Xeon 2.3GHz, 1GB memory). The CMSIM step is therefore mostly CPU bound, even if a non-negligible I/O is needed (~250 MB). The output data of each job was stored in the close Storage Element area and registered into the RLS.
The jobs were submitted in bunches of 30-40 jobs from the CNAF User Interface using both the CERN and CNAF Resource Brokers. The Resource Broker sent jobs to suitable Computing Elements matching the requirements to have the CMS software installed.
The strategy adopted was to explicitly exclude sites once they were showing systematic problems, thus relying on average on about 60-70 CPUs spread over 5-7 sites. The goal was to saturate the available resources with "long" jobs and then test the scalability gradually increasing the number of long jobs in the system up to 400.

More than 600 jobs were submitted: about 200 "fast" jobs and 400 "long" jobs clustered to have about 100 long jobs in the system at the same time, as shown for example in Figure 9.

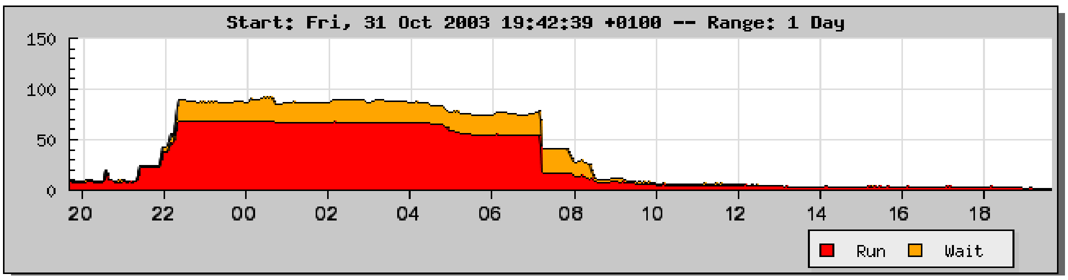

Figure 9: Number of "long" jobs, both running and waiting on LCG-1 system, as a function of time, as shown from GridICE daily graph.

On the basis of limited statistics of submitted jobs the overall efficiency was about 80-90% for "fast" jobs and about 50% for "long" jobs. However, the major sources of troubles encountered during the tests were related to problems at various sites, consisting in both systematic and transient failures.

The systematic failures were due to mis-configuration, for example:

- Jobs failed because the CMS software was not installed or was not found by the job, at least not in all the WNs of a Computing Element, even if the CE advertises in the Information System to have that software.
- The replica manager command failed because of mis-configuration at the site.
- Jobs were aborted at CNAF site with reason *"Got a job held event, reason: Globus error 3: an I/O operation failed"*.
- The application executable failed with segmentation fault in specific Worker Nodes. This behavior could be triggered by lack of disk space but it was not possible to clearly identify the cause.

The observed transient failures, whose cause would have needed more deep investigation, were:

- Jobs aborted because the submission done via CondorG to Globus didn't work ("Globus resource down"). The gatekeeper being unreachable due to CE disconnected because of network failures could cause this.
- Jobs aborted with reason *"Cannot read JobWrapper output, both from Condor and from Maradona"*. The standard output of the script where the user job is wrapped around, that is used to check if the job reached the end successfully, happened to be empty. It is hard to identify the real cause of this problem since there could be many possible reasons.

There were no failures in the file registration into the RLS, also in the case of "fast" jobs (CMKIN), where the probability to have more CMKIN concurrent jobs writing into the Replica Catalog is higher. The problem of too many concurrent jobs writing into the LDAP-based Replica Catalog, causing it to stick, was present last year and it forced to slow down the submission rate for CMKIN jobs.

Very rarely there were problems related to the information system, so few jobs were aborted due to lack of matching resources because the Resource Broker relies in the information service to discover available resources *("Cannot plan: BrokerHelper: Problems querying the information service")*.

The Resource Broker was rather stable, it only happened once that it was hanging and not responding but it was quickly fixed and in any case, there were others Resource Brokers properly working.

4.3 Test and results of OSCAR jobs

The LCG-1-"south testbed" (that includes Italian and Spanish sites) was used in order to run a real OSCAR production as a part of the Pre-Challeng-Production effort .

OSCAR is a "Geant4" [23] based CMS software for simulation of tracking into detectors. It uses POOL [24]as underlying persistency layer . OSCAR requires as input an ntuple, produced by a corresponding CMKIN job and other additional file like a card file (.orcarc) and POOL catalogues. The output is a collection of several EVD files in POOL format that contain the events and their related metadata. They will be the input of the following digitization step in the CMS Simulation chain.

The OSCAR production of a total of 600 K events for dataset "*mu03_tt2mu*" was assigned to run within the LCG-1-"south testbed". Table 5 summarizes the computational characteristics of these jobs.

The input ntuples were stored into the SRB system in CIEMAT (Spain). In order to make them available on the Grid, they were copied into the Padova storage element and registered into the CMS RLS.

The output data upload to an SE and their registration into RLS was not performed with the replica manager copyAndRegisterFile command. The files were copied to the SE with globus-url-copy and, in order to maintain the POOL unique GUID, the POOL FCpublish tool was used to register the produced data into RLS.

	OSCAR			
Dataset	Time*/event	Size/event	Time*/job	Size/job
mu03_tt2mu	~180 s	~0.6MB	~12 hours	~160 MB

Table 5 – Size of data sample and CPU time per event simulation of OSCAR production steps (* *Xeon 2.4 GHz, 2GB*), for the dataset that was simulated within the CMS/LCG-1-south testbed. The CPU time and size per job, each job processing 250 events, is also reported.

The Grid services used during the OSCAR production are listed in Table 6.

Two UIs were used: one in Bari, where the job submission was done using McRunjob, and the other in Padova where CMSProd was used.

While CMSProd jobs were capable of recovering the input ntuple also from SE outside the LCG-1-south testbed boundary, the version of McRunjob used in Bari had some limitation due to local configuration problems. For this reason jobs launched from the Padova UI could run on many LCG-1 sites (Torino, Milano, Legnaro, CNAF, etc.), while the jobs launched from the Bari UI had to run on only some of the sites of the INFN Production testbed

EDG Service	**Used by**	**Site**	**Name**
Resource Brokers	Padova, Bari UI	CNAF Padova Catania	edt003.cnaf.infn.it, grid031.pd.infn.it, grid014.ct.infn.it
Top MDS	Information System	CNAF	edt001.cnaf.infn.it
DBII	all RBs	CNAF Padova	ibm140.cnaf.infn.it, grid041.pd.infn.it
Virtual Organization server	all users	CERN	??
Replica Location Service	Jobs on WN and UI	CERN IT	rlscms.cern.ch

Table 6 –LCG Services location and names.

The job submission started on December 18, 2003, ran over the Christmas holidays and finished on January 20, 2004. The chronological summary of the integrated total amount of OSCAR events produced is shown in Figure 10. The plateau around Christmas holidays is due to the testbed being unattended and to the granularity of the plot extracted from RefDB, where completed assignments (100K events each) are reported.

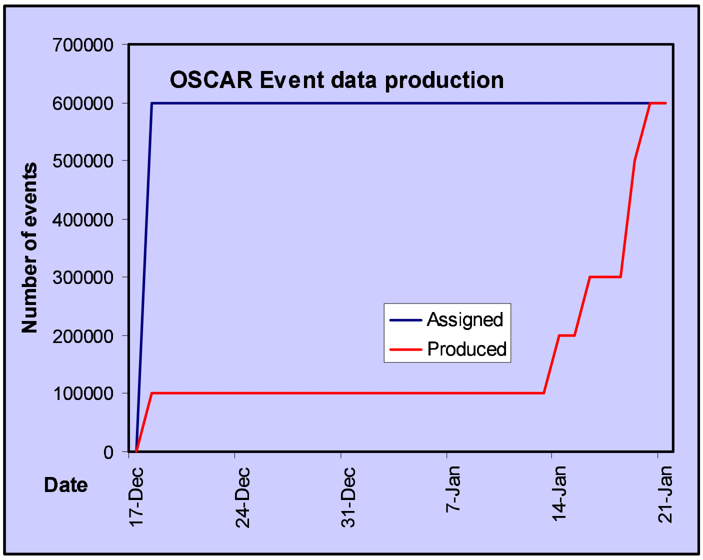

Figure 10: Chronological summary of the integrated amount of OSCAR events produced (from RefDB information, with 100k events resolution) within the LCG-1 south-testbed.

Table 7 summarizes the Grid efficiency of the submitted jobs: about 75% of the jobs successfully terminated for the Grid evaluation, thus reaching the "*Done*" status. Details about of the "*Abort"* status reason are also reported. However the criteria of success and/or failure from CMS point of view are more stringent with respect to a Grid evaluation. In fact a job that reaches the Grid status "*Done*" may fail, from the CMS production point of view, mainly because the access to input data or the staging of the output files was not successful but also due to OSCAR code errors.. The overall percentage of successful jobs from CMS view was of the order of 50-60%. Occasionally there were jobs which successfully completed all steps required by CMS but did not reach the Grid Status "*Done*". All failed jobs were resubmitted (one o more times) in order to complete the entire production assignments.

The main causes of failing jobs could be traced in

- Failure of the central services (services not reachable, RB hardware and middleware software crashes). During the production the resource broker was changed due to some problems with Workload-Manager that leaved bunches of jobs in waiting status.
- Site failures (sites not reachable, PBS_SCHED daemon crashes, SE mis-configurations, losses of Grid services, incorrect queues settings)
- Failures due to user errors (input files not registered in RLS, input file path incorrectly specified, job requirements incorrectly specified)

It was also noted that the information about the job status on the Grid (that is the result of edg-job-status command) was not well synchronized with the real status of job.

In order to compare the LCG-1 results with the LCG-0 ones, it should be pointed out that, in the case of the LCG-0, the farms were set up and maintained by motivated CMS personnel, in the case of LCG-1 these task were taken over by professional operators but in general with little experiences with CMS and GRID software. This led to poorer farm maintenance: as matter of fact the testbed was completely unattended during the Christmas holidays. In this respect it is interesting to note that the GRID performance in this period degraded by about 30%.

Grid job status efficiency for OSCAR jobs	
~ 75%	**Done**
~ 20%	**Aborted** ▪ ~ 43% jobs aborted with reason "*Cannot read JobWrapper output, both from Condor and from Maradona*" ▪ ~ 25% jobs aborted downloading/uploading the sandboxes ▪ ~ 5% jobs aborted with Globus errors like "*Globus resources down*" or "*Got a job held event, reason: Globus error 129 the standard output/error size is different*" or "*Globus error 7: authentication failed*" ▪ ~ 25% jobs aborted in the Resource Broker matchmaking the user required resources (*"Cannot plan: BrokerHelper: no compatible resources"*) ▪ ~ 2% jobs aborted due to information system instability (*"Cannot plan: BrokerHelper: Problems querying the information service")*
~ 5%	**Other status (running, scheduled, waiting…)**

Table 7 – Grid job efficiency expressed in term of grid job status. A detailed breakdown of the failures is also reported.

5 Conclusions

While DataGrid approaching has ended, CMS will continue to use and develop middleware-aware applications. Many key issues of distributed (Grid) computing were addressed during the three years of the DataGrid project and the three years of CMS Computing growth to its final LHC startup design.

There is still a long way behind the goal of CMS computing and also of Grid computing, however a big step forward was achieved.

CMS needs to continue to test new functionalities and organization of services, as LCG-1 (and the coming LCG-2) system is providing. DataGrid gave the start of this process for LHC computing and for many other sciences.

Good efficiencies and stable conditions of the system were obtained in comparison with what obtained in previous challenges [19], showing the maturity of the middleware and of the services, provided that a continuous and rapid maintenance is guaranteed by the middleware providers and by the involved site administrators. The CMS production (and early analysis) tools could be integrated with the Grid services, with good performances.

Many suggestions both for the improvement of Grid services and CMS production/analysis tools were also obtained. Some weak point and font of failures were identified, and whenever possible corrected, demonstrating that the *distributed-grid-model* is approaching a level of usability near to what is required by LHC Experiments.

Apart from the many successes of middleware functionalities and implementation of services, on which CMS built upon its successful productions and "stress tests", some key issues need to be addressed as a priority:

- Stability of services and site configuration consistence is still at a fragile status, even if usable at the price of careful choices and deep knowledge by the users;
- Information system robustness and distribution, including scalability to many (eventually temporarily disconnected) sites;
- Scaling of distributed data (and meta-data) services with a better performance and user available interface (both GUI and APIs);
- Organization of the distributed services and resources (including duplication of them) at different sites that allow a fast and more efficient tracking of unavailability or mis-configuration (even if transients);
- Improved methods for experiments' reservation of resources within the Grid environment, also allowing for customized installations;
- Better definition of the "interfaces" of the functionalities provided by the middleware and the responsibilities (wanted duties) of the applications.

As a conclusion, it should be stressed again that the major results and progresses were obtained when the effort of development and testing was carried on in close contact with the experiment's activities, both for middleware progress and for experiment success. Joint task forces proved to be the key choice.

Acknowledgments

We would like to acknowledge the invaluable help and support of the LCG-EIS and CMS Production Teams.

Without the friendly collaboration and prompt effort of the EDG and EDT people the present work would have not been possible.

Particular thanks are due to all the site managers of the participating sites, and also of some sites not involved, for their assistance and continuous excellent work.

Many colleagues gave their constant help and advice and in particular we would like to thank Andrea Chierici (INFN-CNAF), Nikolai Smirnov (INFN-Padova), David Rebatto (EDT-Milan), Luca Vaccarossa (EDT-Milan) and Tony Wildish (CMS, CERN-Princeton).

References

[1] CMS Experiment: http://cmsdoc.cern.ch/cms/outreach/html/index.shtml

[2] LHC Project: http://lhc.web.cern.ch/lhc/general/gen_info.htm

[3] CERN (Geneva, CH): http://public.web.cern.ch/public/

[4] LCG Project: http://lcg.cern.ch/

[5] EDG Project: http://eu-datagrid.web.cern.ch/eu-datagrid/, EDG Testbed http://marianne.in2p3.fr/datagrid/giis/giis.html

[6] EDT Project: http://datatag.web.cern.ch/datatag/, and GRIDice Project: http://server11.infn.it/gridice/

[7] OCTOPUS, the CMS software repository for "production": http://cmsdoc.cern.ch/swdev/viewcvs/viewcvs.cgi/?cvsroot=OCTOPUS

[8] CMS Internal Note 2002/044, V. Lefebure and J. Andreeva, "RefDB"

[9] G.Graham "McRunJob: A Workflow Planner for Grid production processing" , CHEP03 , S.Diego, California, March 2003.

[10] CMSProd Web page: http://nsmirnov.home.cern.ch/nsmirnov/iam/cms/cmsprod.html

[11] "A Simple Mass Storage System for the SRB Data Grid", M.Wan, A.Rajasekar, R.Moore, P. Andrew, 20th IEEE/ 11th NASA Goddard Conference on Mass Storage Systems & Technologies (MSST2003) San Diego, California, April 7-10, 2003.

[12] McRunjob/EDG plug-in Web page: http://home.cern.ch/fanfani/grid/LCGLite

[13] C.Grandi, A.Renzi "Object Based System for Batch Job Submission and Monitoring (BOSS)", CMS NOTE-2003/005
BOSS Home Page, http://www.bo.infn.it/cms/computing/BOSS/

[14] Virtual Data Toolkit Home Page, http://www.lsc-group.phys.uwm.edu/vdt/

[15] VOMS Home Page, http://hep-project-grid-scg.web.cern.ch/hep-project-grid-scg/voms.html

[16] GLUE schema, http://www.cnaf.infn.it/~sergio/datatag/glue/index.htm

[17] SDSC Storage Resource Broker Home Page, http://www.npaci.edu/DICE/SRB/

[18] Boss2root Web page: http://home.cern.ch/bonacors/boss2root.htm

[19] "CMS Test of the European DataGrid Testbed", CMS and EDG Stress-Test Task Force; CMS NOTE-2003/014.

[20] Savannah Web Page: http://savannah.cern.ch/

[21] Grid.it Web Page: http://www.grid.it/

[22] DataTAG collaboration, S. Andreozzi et al. “Pilot production grid infrastructure for high energy physics applications”, submitted to “Future Generation Computer System (FGCS), Elsevier Pubblishing; special issue of “High-Speed Networks and Services for Data-Intensive Grids”

[23] Geant4 Project: http://wwwasd.web.cern.ch/wwwasd/geant4/geant4.html

[24] POOL Project: http://lcgapp.cern.ch/project/persist/