

<p><b>CMS Data Challenge 04</b>  <b>Event Simulation, Distribution and Analysis.</b></p>	 <p>www.uscms.org</p>	 <p>www.ppdg.net          www.ivdgl.org          www.griphyn.org</p>
--	--	---

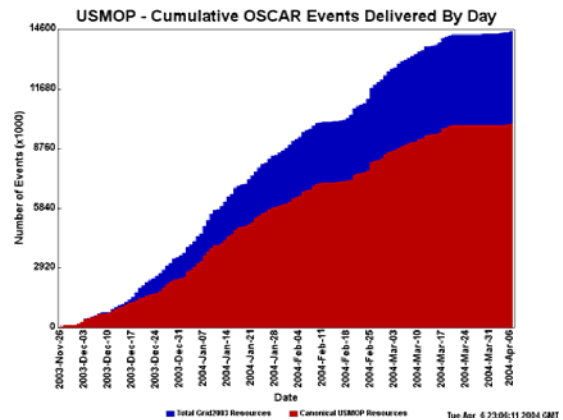
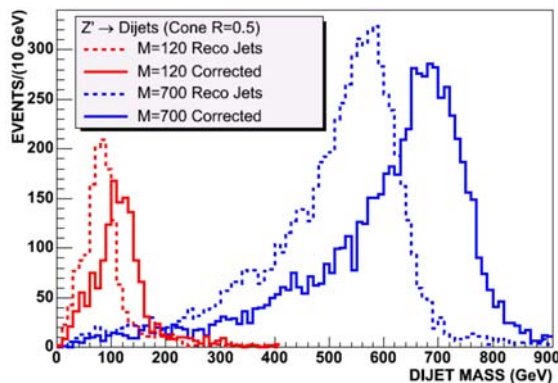
The CMS Data Challenge DC04 is one of the milestones of the experiment scoped to ensure the experiment is ready with its global data distribution and analysis systems for the start of data taking at the Large Hadron Collider at CERN in 2007. DC04 had the full attention of the software and computing program from November 2003 through its completion on May 1 2004.

The performance metrics for DC04 were to provide a baseline to give the experiment input to the Physics and Computing Technical Design Reports in the next two years. These design reports will form the baseline to which the production data processing and analysis systems will be built and must perform.

The goal of DC04 was to perform at 25% of the throughput needed at the start of data taking in 2007 (5% of the full LHC rate) including distribution of data from CERN to the Tier-1 regional centers for 1 month. Additional goals were to provide the software infrastructure to manage and distribute the data, and to provide an end to end demonstration of event reconstruction and analysis to show the state of readiness of the experiment's software infrastructure. The goal was to demonstrate the full chain of data processing and analysis to prove the experiment's software system. It was not overall a challenge to show maximal computational throughput.

The goals of the data challenge were met at some time during the month but were not sustained. The throughput was sustained for 1 day. The final result of analysis of the generated CMS data, distributed from CERN to the worldwide CMS collaboration regional computing centers, was indeed achieved

Pre-Challenge event production was required prior to DC04 in order to produce the simulated events on which the challenge would be run. This production was done using traditional farms, as well as the LCG grid infrastructure in Europe and the Grid3 grid environment in the US. 70M Monte Carlo events were produced, about 25% of them in the US. Significant benefit was obtained through opportunistic use of Grid3 resources – an overall increase in throughput of 50%.



CMS simulation production jobs were submitted to Grid3 using the components MOP and McRunjob that were developed by USCMS and are part of the CMS production environment. The resulting simulated data were written into the dCache-based mass storage system located at the Fermilab Tier-1 center. Simultaneous usage of CPU resources peaked at 1200 CPUs, controlled by a single FTE.


Using Grid3 was a milestone for CMS computing in reaching a new magnitude in the number of autonomously cooperating computing sites for

production. The efficiency of successful jobs was 65%, less efficient than running on a dedicated cluster, but still high enough to make the grid execution mode of value. A variety of errors were encountered from the experiment infrastructure itself, through disks filling up with the output data, to network problems between the sites. Good news was that the support effort required to sustain job throughput was reduced by a factor of 2 compared to May 2002, (the date of the last PPDG news update on CMS event generation).

The DC04 pre-production included the use in the US of storage services presented to the grid. Data transfer was managed between the Tier-2 sites and the Fermilab Tier-1 mass storage system using the DESY/Fermilab SRM/dCache system.

Once generated events had been stored at CERN they were reconstructed on the local farm. The 25Hz throughput translated to about 2000 jobs a day with 40MB/sec transferred from the mass storage system.

A major focus of DC04 was to provide sustained and effective transfer of data between the CERN Tier-0 and distributed Tier-1 facilities across Europe and the US. This proved a challenging application of grid technologies for data management, with the LCG-2 Replica Location Service, SDSC Storage Resource Broker, DESY/Fermilab SRM/dCache all proving ultimately capable after integration and debugging.

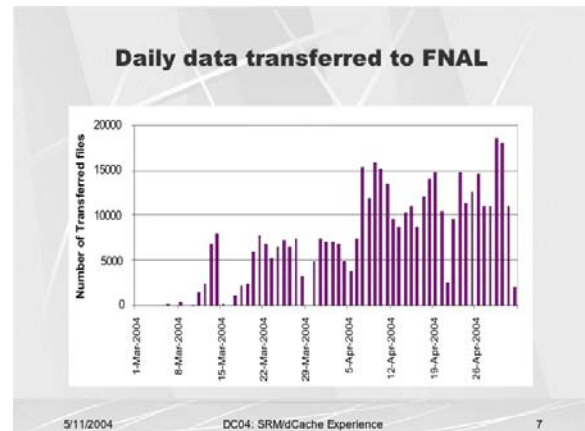


A further challenge was the analysis of the DC04 data. The event data at the remote Tier-1 sites needed to be readable and understood by the 10,000 C++ line experiment analysis framework. This was the first major test of the experiment's data management software.

A typical data sample was that transferred to Fermilab: 5.2 TB of data was transferred in 440K files. The average event size for the data challenge was about 100Kbyte, a factor of 5 less than the final event size for the experiment. As a result the size of

## USCMS Data Challenge 04 is Grid Based

files for distribution was too small to allow efficient management and transfer.



Now the data challenge itself is over one of the most interesting parts of the process starts – analysis of the data by many physicists at the universities and laboratories around the country, for analysis at Fermilab and the prototype Tier-2 centers at the University of Florida, the University of California San Diego and at the California Institute of Technology

The lessons learned from DC04 are being debated and discussed in many CMS meetings. The distributed computing model of the experiment was strengthened by the results of the data challenge.

Besides the need to ramp up the overall throughput by an order of magnitude, the data management system must be able to manage several orders of magnitude more files – or the order of  $10^{*}8$ . Preparations are already starting for DC05 and beyond. Attention will be given to the higher level services such as data management and cataloging needed for the experiments production systems for data taking

The accomplishments of DC04 depended on the combined efforts of the engineers and physicists in the software and computing groups; the CERN, Fermilab and other Tier-1 center central support staff; the support of the Virtual Data Toolkit and other Trillium grid project effort, and the Grid3 operations team.

While December to May was a long five months for many members of CMS, the data challenge was extremely productive and an essential step towards the systems needed to enable new discoveries in 2007.