LHCb Data Challenge 2004

A.Tsaregorodtsev, CPPM, Marseille



Grid Deployment Board, 10 November 2003

Outline

- DC 2004 Goals
- LHCB Software for DC2004
- Resources needed
- DIRAC production system
- Use of the LCG platform
- Conclusions



Data Challenge 2004 – goals

- Main goal: gather information to be used for writing LHCb computing TDR
 - Robustness test of the LHCb software and production system
 - Using software as realistic as possible in terms of performance
 - Test of the LHCb distributed computing model
 - Including distributed analyses
 - Incorporation of the LCG application area software into the LHCb production environment
 - Use of LCG resources as a substantial fraction of the production capacity



DC 2004 – software

 First use of the simulation program Gauss based on Geant4

- Introduction of the new digitisation program, Boole
 With HLTEvent as output
- Robustness of the reconstruction program, Brunel
 - Including any new tuning or other available improvements
 - Not including mis-alignment/calibration
- Pre-selection of events based on physics criteria (DaVinci)
 - AKA "stripping"
 - Performed by production system after the reconstruction
 - Producing multiple DST output streams



DC 2004 – software



Lнср

Grid Deployment Board, 10 November 2003, CERN

Incorporation of the LCG software

- Gaudi will be updated to:
 - Use POOL persistency mechanism
 - Use certain SEAL services
 - e.g. Plug-in manager
- All the applications will use the new Gaudi
 - Should be ~transparent but must be commissioned
- ◆ N.B.:
 - POOL provides existing functionality of ROOT I/O
 - And more: e.g. location independent event collections
 - But incompatible with existing TDR data
 - May need to convert it if we want just one data format



DC 2004 – production scenario

- Generate (Gauss, "SIM" output):
 - 150 Million events minimum bias
 - 50 Million events inclusive b decays
 - + 20 Million exclusive b decays in the channels of interest
- Digitize (Boole, "DIGI" output):
 - All events, apply L0+L1 trigger decision
- Reconstruct (Brunel, "DST" output):
 - Minimum bias and inclusive b decays passing L0 and L1 trigger
 - Entire exclusive b-decay sample
- Store:
 - SIM+DIGI+DST of all reconstructed events



- A total of 47 Million events have been produced in two months
- 18 centres in Europe participated and 80% of CPU outside CERN
- 36 600 jobs have been run and each job:
 producing between 250 and 500 events
 - using from 32 to 56 hours on 1GHz PC
- Produced data stored in ~250 000 files



DC 2003 – summary





Grid Deployment Board, 10 November 2003, CERN

DC 2004 – resources

- CPU requirement is 10 times what was needed for DC'03 – 2.3x10⁹ SI2K hours;
- Current resource estimates indicate DC'04 will last 3 months – 2000-3000 continuously running jobs.
 - Assumes that Gauss is twice slower than SICBMC
 - Currently planned for April-June
- GOAL: use of LCG resources as a substantial fraction of the production capacity
 - We can hope for up to 50%
- Storage requirements:
 - 6TB at CERN for complete DST
 - 19TB distributed among TIER1 for locally produced SIM+DIGI+DST
 - + up to 1TB per TIER1 for pre-selected DSTs

DC 2004 – computing model

Distributed data production

- + As in 2003, will be run on all available production sites
 - Including LCG1
 - Controlled by the production manager at CERN
 - In close collaboration with the LHCb production site managers

Distributed data sets

- + CERN:
 - Complete DST (copied from production centres)
 - Master copies of pre-selections (stripped DST)
- Tier1:
 - Complete replica of pre-selections
 - Master copy of DST produced at associated sites
 - Master (unique!) copy of SIM+DIGI produced at associated sites
- Distributed analysis
 - In limited number of sites.





Distributed MC production system for LHCb

- Production tasks definition and steering;
- Software installation on production sites ;
- Job scheduling and monitoring ;
- Data transfers ;
- Automates most of the production tasks
 - minimum participation of local production managers
- PULL rather than PUSH concept for tasks scheduling
 - Different from the DataGRID architecture.



DIRAC architecture





DIRAC 2.0 Workload Management architecture





Grid Deployment Board, 10 November 2003, CERN

DIRAC use of computing resources

DIRAC design goals to facilitate operation in various environments:

- Scheduling jobs to "any grid" computing resources:
 - "native" sites, running DIRAC Agents;
 - EDG/LCG grid as a whole, passing through RB;
 - EDG/LCG CE's and SE's as DIRAC resources.
- + Flexibility in software installation:
 - Software preinstalled by grid site managers (rpms, etc);
 - Software preinstalled on the sites by the LHCb managers;
 - Software installation on the flight by a running job
- + Using "any grid" storage resources:
 - Data Management component to be able to replicate data between LCG and DIRAC storage elements;



DIRAC status

- DIRAC was successfully used at large scale for the DC 2003:
 * 800-1000 concurrent jobs during 50 days;
- DIRAC is being used in the routine data production for the current LHCb needs:
 - 200-300 concurrent jobs at 4-5 sites;
 - + "0.5 FTE" production manager at CERN.
- The functionality is being upgraded to meet the requirements of the DC2004:
 - Support for analysis tasks;
 - Data Management;
 - User Interface (GANGA);
- The DIRAC architecture review is in progress:
 - Refactoring along the lines of the ARDA architecture proposal
 - Use ARDA FileCatalog, DataManagement, IS, ...
 - POOL ARDA FileCatalog compatibility to be solved.



Tests of the LCG platform

- Tests on the EDG1.4 testbed (Feb 2003):
 - Standard LHCb production tasks;
 - ~35% success rate.
- EDG2.0/LCG-1 tests are in progress:
 - Automatic packaging of the LHCb software suitable for installation on the LCG testbed is ready;
 - LHCb production tasks run on the EDG 2.0 platform with their output stored on the EDG SE's;
 - Submission through the DIRAC WMS;
 - Submission to EDG RB or directly to CE's.
 - Analysis tasks using previously stored input data are in preparation
 - Large scalability tests to follow.
- There are still developments to be done but by the time of DC2004 the LCG platform will be one of the fully incorporated LHCb "production sites"



Use of LCG resources in DC 2004

- We hope to use up to 50% of the LCG resources for the needs of the DC 2004:
- This is conditioned by:
 - The availability of major non LHC dedicated centers both through usual and LCG workload management system;
 - The LCG Data Management access to major MSS (CERN, IN2P3,FZK,CNAF,RAL,Spain);
 - The overall stability and efficiency (>95%) of the system.



- Data Challenge 2004 is essential for the Computing TDR preparation;
- DC 2004 will test the robustness of the LHCb software and computing model;
- LCG software components will be incorporated into the DC2004 software;
- DC 2004 will make use of the LCG resources as much as there will be available.

