



Enabling Grids for
E-science in Europe

www.eu-egee.org

*EGEE-2 Conference, Geneva
25th-29th September*

Report on the LHCb DC06

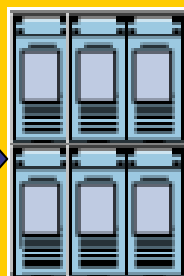
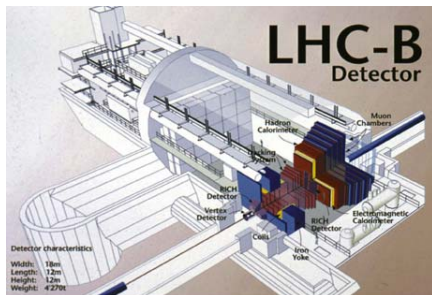


EGEE is a project funded by the European Union under contract IST-2003-508833

DC06 aims

- LHCb DC06 goals
 - to produce simulated data for the "physics book"
 - to make a realistic test of the computing model in order to mimic what LHCb will have to do with real data
- DC06 is composed by different (concurrent) steps
 - MC Simulation
 - Reconstruction
 - Stripping
 - Data Transfer (data upload, redistribution and replication)

} sc4
- Analysis is not part of DC06!



Farm Online



**DST+RAW+TAG
~12 MB/s each T1**

NIKHEF/SARA



RAW 60MB/s



RAI

RAW-mc

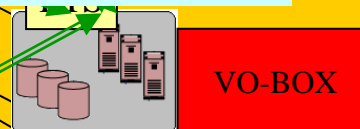
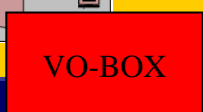
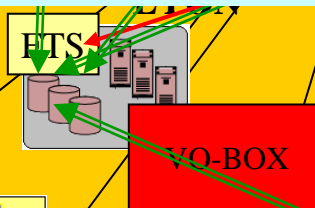
VO-BOX

GRIDKA



VO-BOX

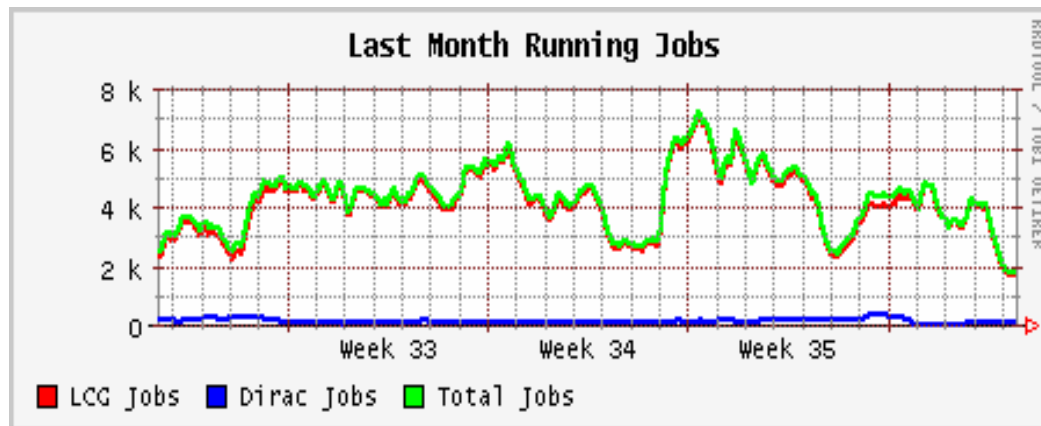
Simulation throughput:
 5K/jobs,
 450MB each job → 20-30MB/s T1
 DST+RAW+TAG redistributed over all T1s



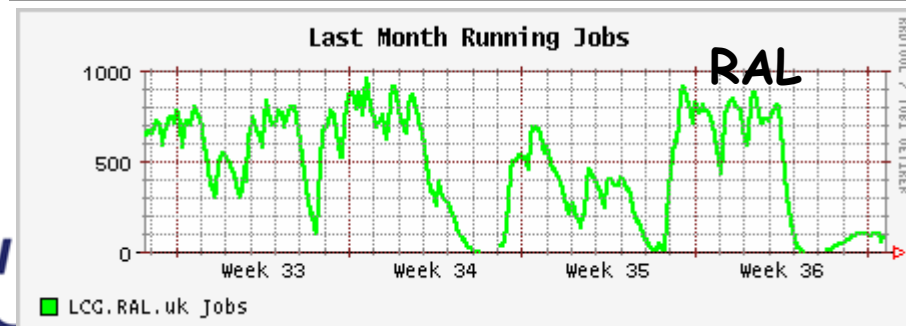
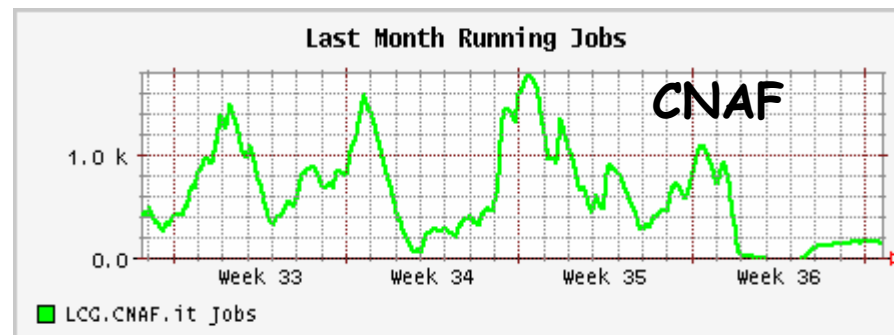
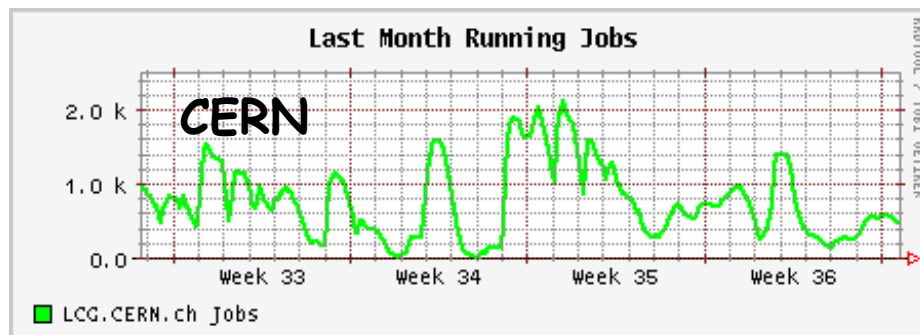
1st step: MC simulation

- The DC06 simulation phase started at the beginning of May
- The Physics quality simulation started at the beginning of July
 - Physics goal is to simulate (with Gauss v24r8 onward)
 - 25M bbar with high luminosity ($5 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$)
 - 25M bbar with low luminosity ($2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$)
- DIGI (simulated raw data + MC truth) produced over all sites (T0, T1s & T2s) and stored only at CERN Castor SE

Running jobs



- ~5k running jobs on average with a peak at 7k during the last month. Reconstruction and simulation jobs had often to compete them self:
prioritization mechanism was/is highly desirable

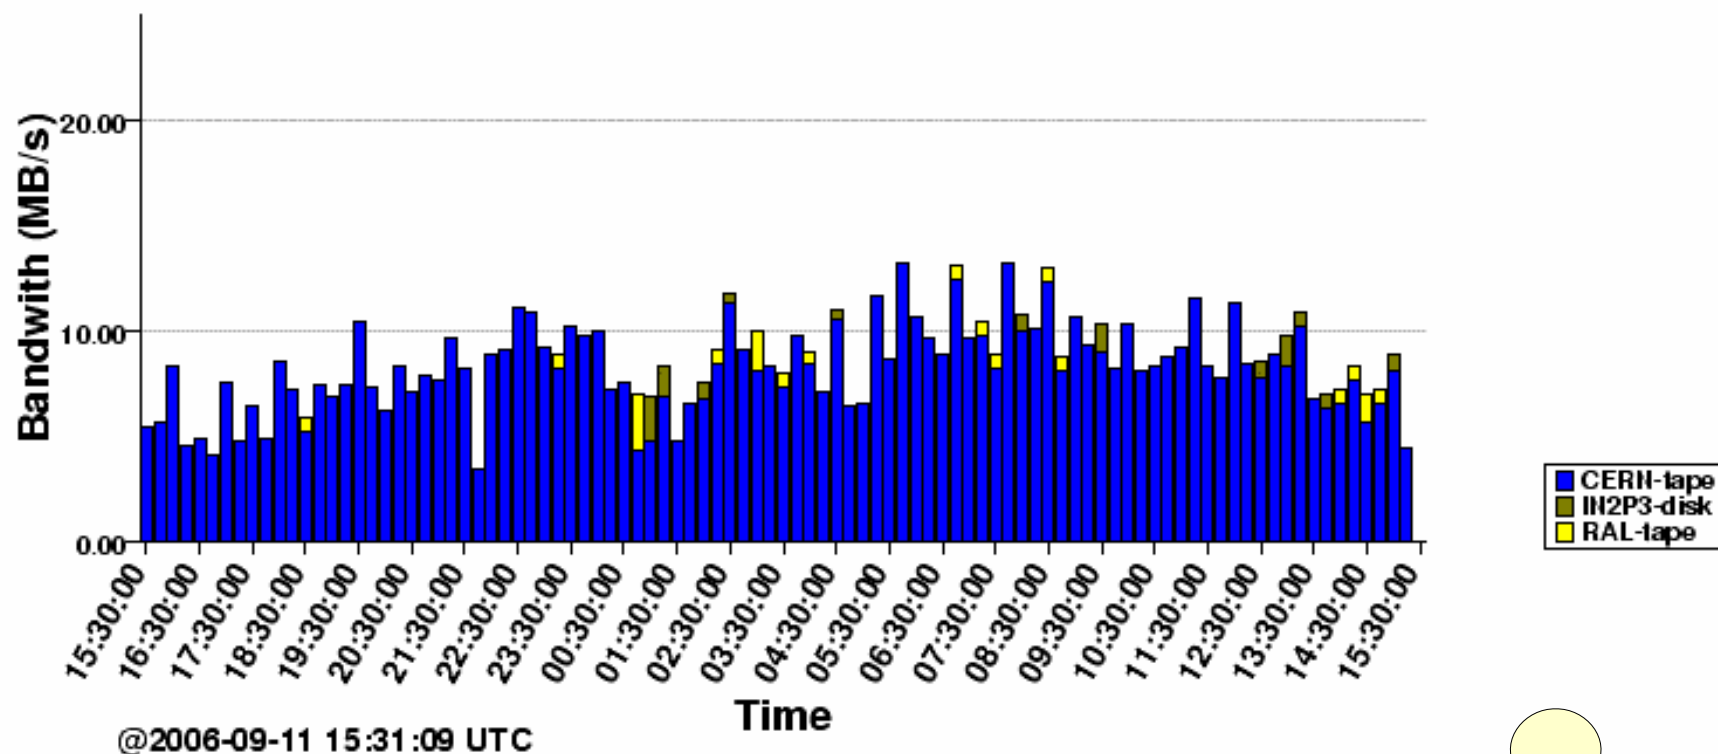


- The periodic decrease of running job is due to the fair-share on the site
 - Site shouldn't brutally drop to zero the share but they should rather reduce it gradually



Simulation data by destination

24 Hours Bandwidth for srm protocol by Destination



- If CERN is not available, the other large sites (T1) are used as T0.

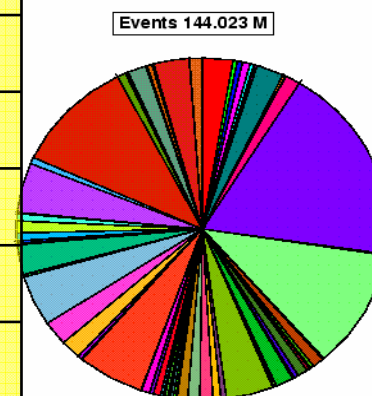
- Reliability and redundancy of central (critical) services is the key-word
- Fail-over mechanism and geographical distribution of data next-to-do

Simulation accounting since May 2006

Breakdown of sites

Site	Events (%)	Events
LCG.CERN.ch	18.59	26,773,752
LCG.CNAF.it	11.18	16,105,760
LCG.RAL.uk	10.60	15,264,775
LCG.LPC.fr	5.81	8,372,270
LCG.Manchester.uk	4.76	6,853,250
LCG.QMUL.uk	4.69	6,754,999
LCG.GRIDKA.de	4.10	5,910,519
LCG.USC.es	3.07	4,420,500
LCG.NIKHEF.nl	2.82	4,055,101
DIRAC.Lyon.fr	2.64	3,807,749
LCG.Barcelona.es	2.37	3,417,810
LCG.Liverpool.uk	2.27	3,269,500

Total number of events produced 144M over ~100 sites



DIRAC.Lyon.fr	2.64%
DIRAC.Zurich-MH.ch	0.38%
DIRAC.Zurich.ch	0.46%
DIRAC.joel.ch	0.00%
LCG.ACAD.bg	0.67%
LCG.AUVER.fr	0.32%
LCG.BHAM-HEP.uk	0.10%
LCG.BIFI.es	0.20%
LCG.Barcelona.es	2.37%
LCG.Bari.it	0.23%
LCG.Bologna.it	0.03%
LCG.Bristol.uk	0.01%
LCG.Brunel.uk	1.34%
LCG.CERN.ch	18.59%
LCG.CESGA.es	0.01%
LCG.CGS.fr	0.05%
LCG.CNAF-GRIDIT.it	0.07%
LCG.CNAF.it	11.18%
LCG.CNB.es	0.14%
LCG.CPPM.fr	1.00%
LCG.CSCS.ch	0.32%
LCG.CY01.cy	0.09%
LCG.Cagliari.it	0.00%
LCG.Cambridge.uk	0.00%
LCG.Catania-fails.it	0.01%
LCG.Catania.it	0.26%
LCG.Dortmund.de	0.89%
LCG.Durham.uk	0.41%
LCG.EELA-CIEMAT.es	0.09%
LCG.EELA-UFRJ.br	0.00%
LCG.ELTE.hu	0.04%
LCG.Edinburgh.uk	0.03%
LCG.FORTH.gr	1.54%
LCG.Ferrara.it	0.01%
LCG.Firenze.it	0.20%
LCG.GR-01.gr	0.07%
LCG.GR-04.gr	0.04%
LCG.GR-05.gr	0.01%
LCG.GRIDKA.de	4.10%
LCG.GRNET.gr	0.47%
LCG.Glasgow.uk	0.69%
LCG.HG-02.gr	1.13%
LCG.HG-04.gr	1.05%
LCG.HG-06.gr	0.92%
LCG.HPC2N.se	0.00%
LCG.ICI.ro	0.11%
LCG.IHEP.su	0.21%
LCG.IN2P3.fr	0.26%
LCG.INR.ru	0.10%
LCG.IPP.bg	0.08%
LCG.IPSI-IPGP.fr	0.04%
LCG.ITEP.ru	0.27%
LCG.ITWM.de	0.01%
LCG.Imperial.uk	0.06%
LCG.JINR.ru	0.05%
LCG.KFKI.hu	0.23%
LCG.KIAE.ru	0.63%
LCG.KIAM.ru	0.31%
LCG.KNU.kr	0.02%
LCG.Krakow.pl	0.72%
LCG.LAL.fr	0.05%
LCG.LPC-fails.fr	0.01%
LCG.LPC.fr	5.81%
LCG.LPN.fr	0.11%
LCG.Lancashire.uk	0.35%
LCG.LesC.uk	0.00%
LCG.Legnaro.it	1.79%
LCG.Liverpool.uk	2.27%
LCG.Manchester.uk	4.76%
LCG.Milano.it	0.20%
LCG.Montreal.ca	0.00%
LCG.NCP.pk	0.02%
LCG.NCU.tw	0.01%
LCG.NIKHEF.nl	2.82%
LCG.Napoli-Atlas.it	0.15%
LCG.Napoli.it	0.22%
LCG.OU.il	0.02%
LCG.Oxford.uk	0.73%
LCG.PAKGRID.pk	0.00%
LCG.PIC.es	1.11%
LCG.FNPI.ru	0.02%
LCG.Padova.it	0.67%
LCG.Pisa.it	0.12%
LCG.QMUL.uk	4.69%
LCG.RAL-HEP.uk	0.61%
LCG.RAL.uk	10.60%
LCG.RHUL.uk	0.75%
LCG.SARA.nl	0.07%
LCG.SINP.ru	0.19%
LCG.SRCE.hr	0.00%
LCG.Sheffield.uk	1.60%
LCG.Sofia.bg	0.02%
LCG.TAU.il	0.09%
LCG.TCD.ie	0.02%
LCG.Torino.it	0.42%
LCG.UCL-CCC.uk	0.00%
LCG.ULAKBIM.tr	0.17%
LCG.USC.es	3.07%
LCG.WARSZAWA-failed.pl	0.04%
LCG.WARSZAWA.pl	1.03%
LCG.WCSS.pl	0.00%
LCG.WEIZMANN.il	0.06%

@2006-09-08 Between 2006-05-10 - 2006-09-07

Simulation Accounting since May 2006

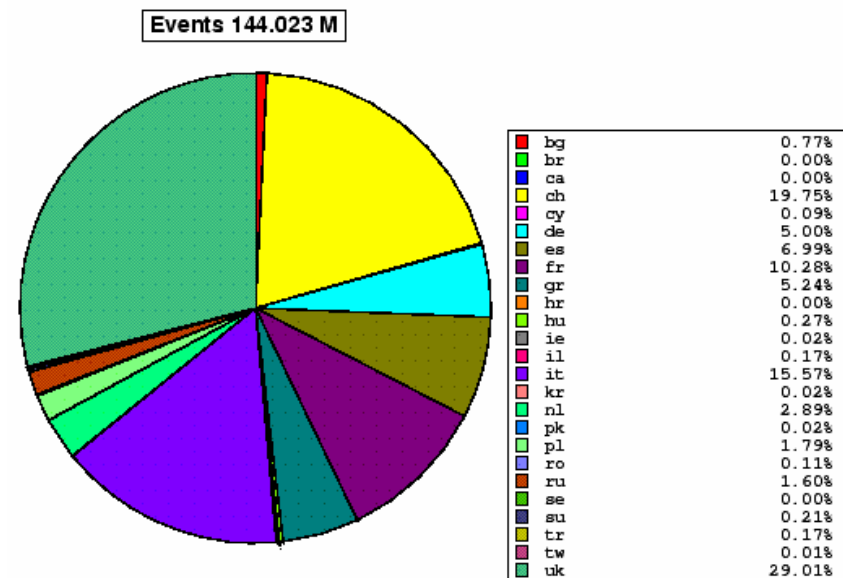
Breakdown of countries

Site	Events (%)	Events
------	------------	--------

Total number of

The number of simulated events requested for physics studies has been fulfilled (July-August)

IT	15.57	22,429,760
FR	10.28	14,799,830
ES	6.99	10,065,229
GR	5.24	7,552,039
DE	5.00	7,206,019
NL	2.89	4,156,350
PL	1.80	2,584,749
RU	1.60	2,305,570
BG	0.77	1,108,510



@2006-09-08 Between 2006-05-10 - 2006-09-07

2nd and 3rd step: reconstruction/stripping

- The DIGI files are distributed over all T1s where the reconstruction will be performed accordingly a variable share via automatic requests.
 - IN2P3, CNAF, PIC, RAL, GRIDKA, CERN, NIKHEF/SARA
- Reconstruction implies direct access of the application (Brunel) to data on the *on the T1's "tape1disk0"-SE*.
- The produced rDST (reduced DST) files are stored where they have been processed ready for subsequent stripping phase.
- Preselection algorithms on the rDST files are run at all T1 (stripping)
- Produced DST and RAW data of selected events + event TAG files are replicated to all the 6T1+ CERN for subsequent analysis.
- Stripping is not running because the Experiment Application is not yet ready for exploiting the foreseen workflow.

rDST Vs DIGI production

Site	Reco jobs	Reco Events	Reco ratio:	Simu Events	Simu ratio
CERN	593	5.93M	25%	26.77M	18.6%
CNAF	14	0.14M	<1%	16.10M	11.2%
GRIDKA	28	0.30M	1%	5.91M	4.1%
IN2P3	715	7.15M	30%	3.80M	2.6%
NIKHEF/SARA	-	-	-	4.05M	2.9%
PIC	451	4.51M	19%	3.41M	2.4%
RAL	563	5.63M	24%	15.26M	10.6%
TOTAL		23.6M	100%	75.3M	52.4%

- Never run reconstruction with all the sites simultaneously
 - *Data access increases the number of difficulties*
 - foresee special SFT/SAM tests also for accessing files*

Conclusion

- DC06 has been running successfully now for 3 months
 - Most components (excepting stripping) extensively tested on the grid
 - Starting from the file registration when a data will come from the pit all chain to reconstruct events is completely automatic and well tested
 - Only stripping phase is missed but the mechanism to produce stripped data is the same of the reconstruction. No particular problem will be expected
- Simulation
 - runs smoothly since months
 - The first 50M bb events have been produced
 - we are to start the biased sample production
- Reconstruction
 - ~24M of events reconstructed at CERN(25%), RAL(24%), PIC(19%), IN2P3(30%).
 - **Never run over all T1's sites at the same time**
- Stripping
 - The rDST files are ready to be processed: few modification required on Experiment Application to get it running.
- Data Transfer
 - No evident problem related to the applications used to transfer data or outages of central services (either LFC or FTS)
 - The infrastructure has been successfully exercised up to the 50% of the expected final rate
 - The failures are mainly related to problems at the sites
 - **SRM end point, Stagers, power supply failure, gridFTP servers overloaded, Network configuration**