


LHCb Grid Experience

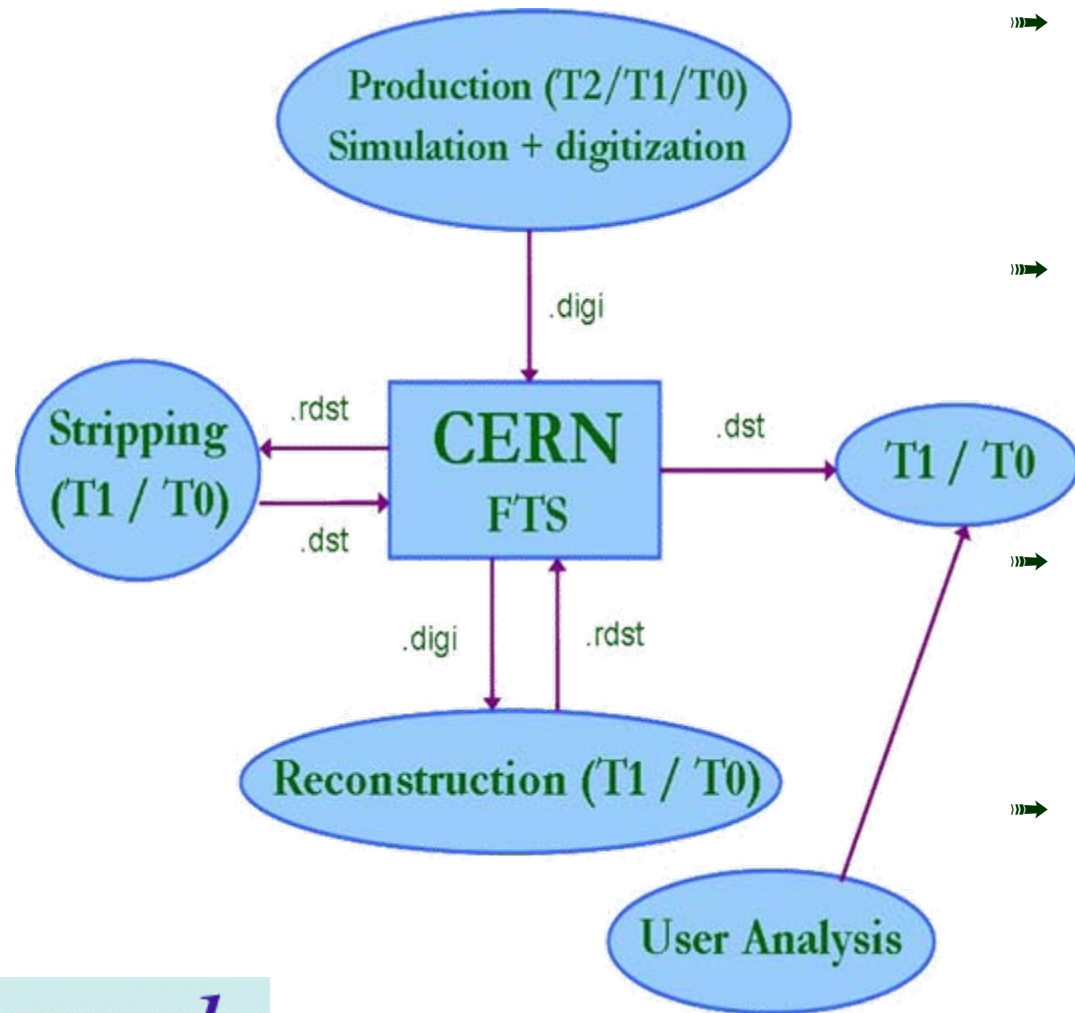
Raja Nandakumar



And many others who helped me with this talk



LHCb computing model



- » CERN (Tier-0) is the hub of all activity
 - Full copy at CERN of all raw data and dst-s
 - All T1s have a full copy of dst-s
- » Simulation at all possible sites (CERN, T1, T2)
 - LHCb has used about 120 sites on 5 continents so far
 - Donot need storage resources on T2
 - Shared area, cpu
- » Reconstruction, Stripping and Analysis at T0 / T1 sites only
 - Most of LHCb storage is here
 - Some analysis may be possible at "large" T2 sites in the future
- » Almost all the computing (except for development / tests) will be run on the grid.
 - Large productions : production team
 - Ganga (Dirac) grid user interface

»»» DIRAC is LHCb's interface to the grid


- Written mostly in python
- Pilot agent paradigm
 - Fine grained visibility of grid to the jobs and DIRAC3 servers

»»» DIRAC3

- Re-writing of DIRAC using 4+ years of experience
- Main ideas / framework retained
 - Many changes in algorithms and implementations
- Security : Authentication & logging for all operations
- Separate out generic and LHCb-specific modules
- Better designed to support more options
 - srm v2, gLite WMS, generic pilots, ...
 - Job throttling, job prioritisation, generic pilots, ...

nraja@lhcb_user (/C=UK/O=eScience/OU=CLRC/L=RAL/CN=raja.nandakumar)

DIRAC



Info	Systems	Jobs	Data	Selected Setup LHCb-Development		
jobs > Monitor						

Site: LCG.RAL.uk	Monitor	Job ID: <input type="text"/>	Sort by: JobID Descending	
Job Status: All	Production monitor Site summary Pilot summary	Updated after: <input type="text"/>	JobGroup: All	<input type="button" value="Submit"/>

Select: All None Action: Reset Kill Delete 1 2 ... 73 >> 25 jobs, out of: 1812

JobId	Status	MinorStatus	ApplicationStatus	Site	JobName	LastUpdate [UTC]	LastSignOfLife [UTC]
<input type="checkbox"/> 18700	■ Matched	Assigned	unknown	LCG.RAL.uk	00000106_00005718	2008-03-06 08:27:35	2008-03-07 07:41:44
<input type="checkbox"/> 18703	■ Done	Execution Complete	Failed To Save Job Outputs	LCG.RAL.uk	00000106_00005721	2008-03-07 05:59:46	2008-03-07 05:59:46
<input type="checkbox"/> 18756	■ Matched	Assigned	unknown	LCG.RAL.uk	00000106_00005745	2008-03-06 08:27:51	2008-03-07 07:42:52
<input type="checkbox"/> 18757	■ Matched	Assigned	unknown	LCG.RAL.uk	00000106_00005746	2008-03-06 08:51:17	2008-03-07 07:42:19
<input type="checkbox"/> 18758	■ Matched	Assigned	unknown	LCG.RAL.uk	00000106_00005747	2008-03-06 09:00:07	2008-03-07 07:46:04

Job ID: 18703








- JDL
- Attributes
- Parameters
- Logging info
- StandardOutput
- Get LogFile
- Actions

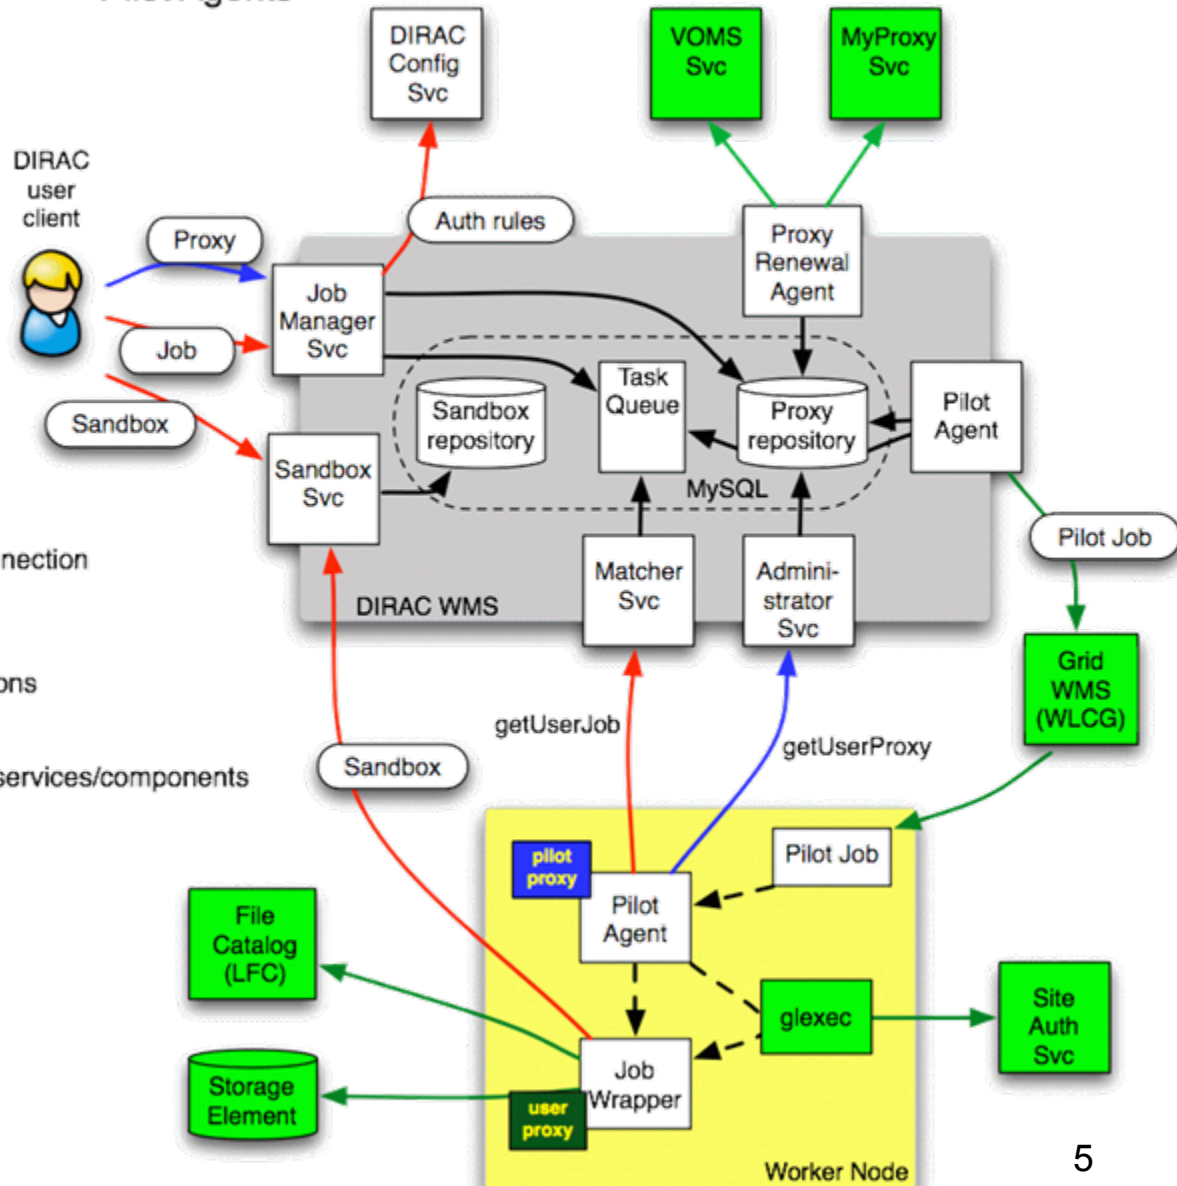
»»» Authentication

- Not needed for browsing the jobs
- Needed to perform actions

DIRAC WMS with generic Pilot Agents

DIRAC3 Workflow

-  Proxy download over DISET secure connection
-  DISET secure connection
-  Local database connection
-  GSI authenticated grid service connections
-  Local job execution (spawning)
-  - Grid services
-  - DIRAC services/components



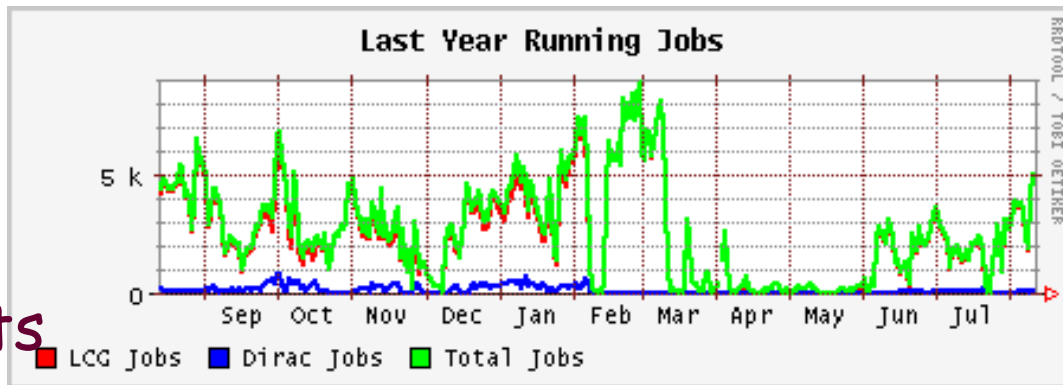


On worker node

- »» Pilot job on the worker node : Regular LCG job
 - Install job agent, check environment ...
- »» Talks to the DIRAC3 WMS
 - Download job with highest priority
 - Gets the required proxy of the job
- »» Implements site policy using glxexec
 - Primarily the site authorisation rules
 - Also extend proxies as needed through glxexec
 - In case of access denial, clean up the workload area and look for next matching job
 - In case of approval, pass job wrapper script + user proxy to the glxexec utility for execution
 - Separate process spawned by job wrapper
 - Job wrapper runs in parallel to monitor the job
- »» When the job finishes
 - Check environment and look for next matching job
 - Clean up and exit when no matching jobs are available
- »» More information :
 - <https://twiki.cern.ch/twiki/bin/view/LHCb/DIRAC3Discussion>
 - <http://indico.cern.ch/getFile.py/access?sessionId=4&resId=2&materialId=0&confId=20230>

»»» DC06 computing challenge

- DC04 before that
 - Data removed now
- Many millions of events simulated, moved, reconstructed




»»» Ran until July 2007

- 120 sites, 5 continents
- Clean up / small productions ongoing
- Use DIRAC2 for all this

»»» CHEP 07 posters, talks

Total Running Jobs: 9715
DIRAC: 0.29% LCG: 99.71%



DIRAC, Zurich-MH, ch	0.14%	LCG, IRB, hr	0.11%
DIRAC, Zurich, ch	0.14%	LCG, ITP, ru	0.21%
LCG, ACAD, bg	0.36%	LCG, IFKI, hu	0.73%
LCG, AUVER, fr	0.24%	LCG, KIAE, ru	0.29%
LCG, Barcelona, es	1.08%	LCG, KIAH, ru	0.09%
LCG, BHM-HEP, uk	0.03%	LCG, Krakow, pl	1.41%
LCG, Bologna, it	0.01%	LCG, LAL, fr	0.84%
LCG, Cagliari, it	0.06%	LCG, Lancashire, uk	2.99%
LCG, Catania, it	0.15%	LCG, LAPP, fr	0.02%
LCG, CERN, ch	25.42%	LCG, LISA, nl	0.63%
LCG, CESSA, es	0.04%	LCG, Liverpool, uk	0.61%
LCG, CGG, fr	0.27%	LCG, LPC, fr	0.75%
LCG, CHAF-GRIDIT, it	0.04%	LCG, LPN-fails, fr	0.10%
LCG, CHAF, it	3.87%	LCG, LPN, fr	0.31%
LCG, CHB, es	0.15%	LCG, Manchester, uk	15.35%
LCG, CPPM, fr	0.69%	LCG, Napoli-Atlas, it	0.02%
LCG, CSCS, ch	0.12%	LCG, Napoli, it	0.03%
LCG, Dortmund, de	0.36%	LCG, IIRHEF, nl	3.32%
LCG, Durham, uk	0.61%	LCG, OU, il	0.04%
LCG, Edinburgh, uk	0.04%	LCG, Oxford, uk	0.54%
LCG, EELA-CIEMAT, es	0.21%	LCG, Padova, it	0.24%
LCG, ETF-RTH, lv	0.14%	LCG, Pisa, it	0.20%
LCG, Ferrara, it	0.19%	LCG, PNP1, ru	0.37%
LCG, FESB, hr	0.08%	LCG, QMUL, uk	10.66%
LCG, FORTH, gr	0.89%	LCG, RAL-HEP, uk	1.61%
LCG, Glasgow, uk	3.63%	LCG, RHUL, uk	0.81%
LCG, GR-01, gr	0.09%	LCG, SIHP, ru	0.34%
LCG, GR-05, gr	0.10%	LCG, Sofia, bg	0.06%
LCG, GRIDKA, de	9.04%	LCG, TCD, ie	0.21%
LCG, HG-02, gr	0.59%	LCG, Torino, it	0.51%
LCG, HG-06, gr	1.16%	LCG, ULAKBIM, tr	0.46%
LCG, HPC211, se	0.39%	LCG, USC, es	1.55%
LCG, Imperial, uk	0.07%	LCG, WARSZAWA, pl	0.47%
LCG, IN2P3, fr	3.31%	LCG, WCSS, pl	0.11%
LCG, IPSL-IPGP, fr	0.06%	LCG, WEIZMANN, il	0.02%

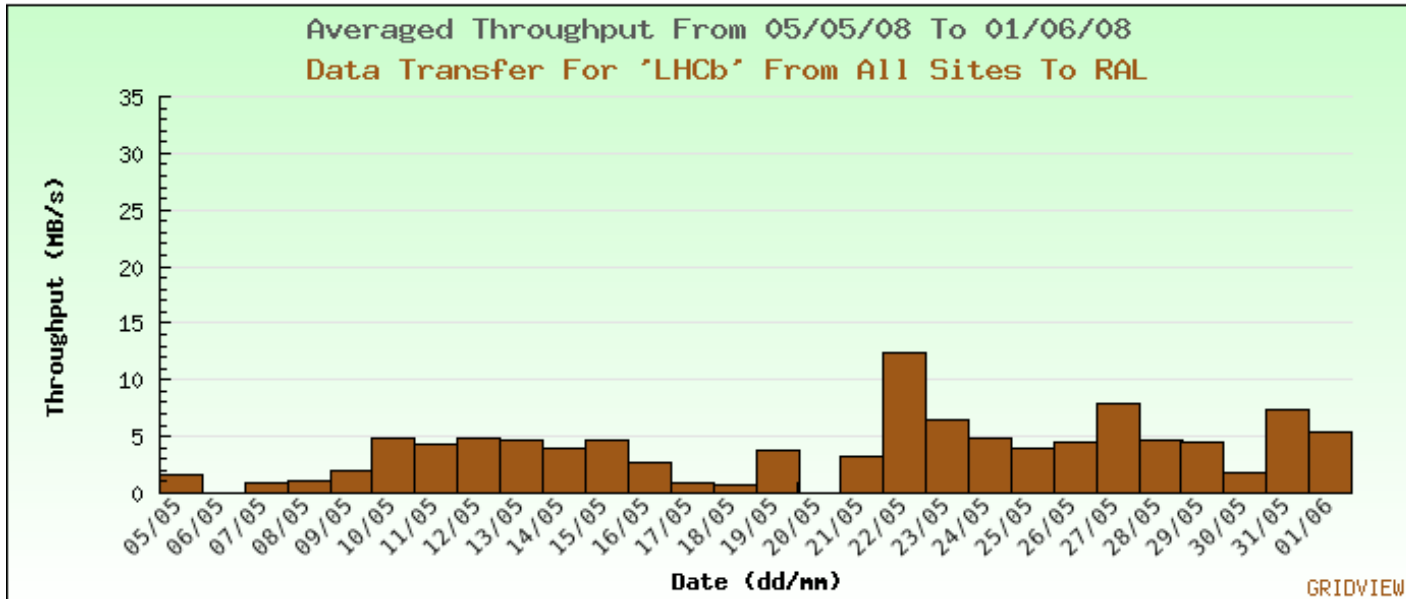
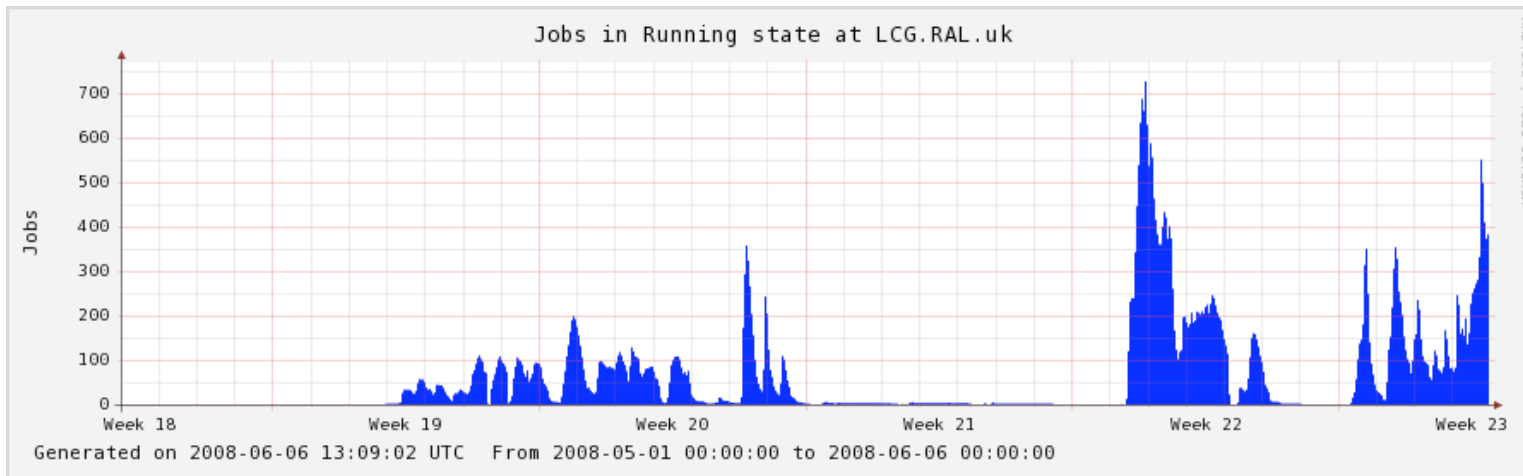
Feb 28 2007, 08:40



Grid Experience – CCRC08

- »» Replicate evolution of data during LHC runs
 - Equivalent of 1 month of running, 50% machine cycle efficiency
 - Data distribution from pit to T0, then T1
 - rfcv : Pit to T0
 - FTS : Out of T0
 - Stable, match or exceed requirement during stable running
 - Reconstruction, stripping at T0, T1
 - Problems at various sites - primarily dCache
 - Open files directly on the SE
 - ⇒ Fall back option of copy to worker nodes possible
 - ⇒ Can lead to other problems however at some sites
 - Use srm2.2 on all sites
 - Testing xrootd protocol now
 - Re-distribution of DSTs between T0, T1 sites using FTS
 - We had problems with our bookkeeping
 - Being intensively developed right now
- »» Improved error reporting in workflow and pilot logs needed for future


May CCRC08 at RAL



- »»» Prepare intensively for data taking
- »»» For T1 centres
 - Restart CCRC type activities once DIRAC3 is ready
 - The stripping is being finished through DIRAC2
- »»» At the T2 centres
 - Test MC production using DIRAC3
 - Run small scale productions
 - Broadly ~1500 jobs a day, 20% in the UK
 - All of the MC production should migrate from DIRAC2 to DIRAC3 over the next few months
 - Small scale DC06 requests are currently being met through DIRAC2
 - Also some user simulation and analysis also running
 - As mentioned, do not need much storage on T2s
 - If we have large enough storage, can use it for analysis
 - ⇒ No formal structures in place for this as yet

LHCb jobs on DIRAC2


Total Running Jobs: 408
DIRAC: 0.00% LCG: 99.51%



GLITE.BMGgrid.hu	0.49%	LCG.NAPOLI-ATLAS.it	0.25%
LCG.BHAM-HEP.uk	0.74%	LCG.Oxford.uk	1.23%
LCG.Bristol-HPC.uk	0.25%	LCG.PIC.es	1.48%
LCG.Bristol.uk	0.49%	LCG.Pisa.it	3.20%
LCG.Brunel.uk	0.49%	LCG.PNPI.ru	0.25%
LCG.Cagliari.it	0.74%	LCG.RAL-HEP.uk	14.29%
LCG.CERN.ch	30.79%	LCG.RAL.uk	0.25%
LCG.CHAF-T2.it	0.25%	LCG.Sheffield.uk	4.93%
LCG.CHAF.it	1.23%	LCG.SRCE.hr	0.74%
LCG.CSCS.ch	0.74%	LCG.TCD.ie	2.96%
LCG.DESY.de	1.48%		
LCG.Dortmund.de	2.71%		
LCG.GRIDKA.de	5.91%		
LCG.HG-02.gr	0.49%		
LCG.IHEP.su	0.74%		
LCG.II-BGU.il	1.72%		
LCG.ITEP.ru	0.99%		
LCG.ITM.de	0.25%		
LCG.JINR.ru	5.42%		
LCG.KIAC.ru	0.25%		
LCG.LAL.fr	4.68%		
LCG.Lancashire.uk	1.48%		
LCG.LAPP.fr	1.23%		
LCG.Legnaro.it	0.74%		
LCG.Liverpool.uk	0.25%		
LCG.LPC.fr	0.25%		
LCG.LPN-fails.fr	1.97%		
LCG.LPN.fr	1.97%		
LCG.Manchester.uk	0.99%		
LCG.Milano.it	0.74%		

Jun 20 2008, 09:16

Total Running Jobs: 1895
DIRAC: 1.90% LCG: 97.68%



DIRAC.Zurich-MH.ch	0.90%	LCG.Lancashire.uk	1.79%
DIRAC.Zurich.ch	1.00%	LCG.LAPP.fr	0.63%
GLITE.BMGgrid.hu	0.42%	LCG.Legnaro.it	2.27%
LCG.Barcelona.es	0.58%	LCG.LPC.fr	1.16%
LCG.Bari.it	0.37%	LCG.LPN-fails.fr	1.95%
LCG.BHAM-HEP.uk	2.16%	LCG.Manchester.uk	3.85%
LCG.Bristol-HPC.uk	0.53%	LCG.NAPOLI-ATLAS.it	0.21%
LCG.Brunel.uk	5.22%	LCG.NIKHEF.nl	2.01%
LCG.Cambridge.uk	1.00%	LCG.Oxford.uk	5.33%
LCG.CERN.ch	21.27%	LCG.PIC.es	3.64%
LCG.CHAF-T2.it	1.95%	LCG.Pisa.it	1.95%
LCG.CHAF.it	2.16%	LCG.PNPI.ru	1.00%
LCG.CPPM.fr	0.79%	LCG.RAL.uk	0.26%
LCG.DESY.de	0.58%	LCG.RHUL.uk	1.74%
LCG.Dortmund.de	1.27%	LCG.ROMA2.it	0.37%
LCG.GRIDKA.de	3.85%	LCG.Sheffield.uk	1.48%
LCG.HellasGrid.gr	0.95%	LCG.SINP.ru	0.05%
LCG.HG-02.gr	0.37%	LCG.SNS-PISA.it	0.32%
LCG.IFH.de	1.16%	LCG.SRCE.hr	0.95%
LCG.IHEP.su	0.53%	LCG.TCD.ie	4.42%
LCG.IN2P3.fr	4.64%	LCG.Trieste.it	3.17%
LCG.INR.ru	0.74%	LCG.USC.es	4.43%
LCG.IPSL-IPGP.fr	0.16%		
LCG.IRE.hr	0.42%		
LCG.ITEP.ru	1.21%		
LCG.ITM.de	2.16%		
LCG.JINR.ru	2.69%		
LCG.KIAC.ru	0.69%		
LCG.KIAC.ru	0.11%		
LCG.LAL.fr	1.11%		

Jun 20 2008, 10:10

To be run on
DIRAC3 soon
(hopefully)