

# Overview of applications view of the job management middleware

WP1-2/Applications joint session  
Heidelberg, September 26, 2003

**Mario Reale**

mario.reale@roma1.infn.it



# Summary



- ◆ EDG and LCG provided GRID demonstrators
  - EDG 2.0.3 application testbed / EDG 2.0 devel testbed
  - LCG-1 certification testbed and prototypal LCG-1 GRID computing service
- ◆ test code sources and on going application activities
- ◆ test results on workload management system
- ◆ most relevant currently opened bugs
- ◆ outstanding general issues/concerns
- ◆ conclusions

# EDG development and application testbeds



- ◆ EDG **2.0 devel** testbed ( Jul-Aug 2003, **2.0.0** tagged on **Aug 28** )
  - EDG dev core sites : **CERN**, RAL, NIKHEF, LAL, CNAF, LYON
  - RB in Lyon, Top R-GMA at RAL
  - Many **temporary development tags** July-Aug ( "t20030822\_1000" )
  
- ◆ EDG **2.0.3 application evaluation** testbed (mid Sept,...growing)
  - Currently **11 sites** ( IFAE,RAL,LYON,NIKHEF,CT,PD,RM,FE,UCL,IC,FZK)
  - RB at NIKHEF, Top R-GMA at RAL

# LCG CERN Certification testbed set up (as of Sept 8, 2003 ) **LCG-CT**



- ◆ 6 CERN-hosted clusters of Linux 7.3 PCs ( 7 different CEs )
- ◆ Deploying **prototypal LCG-1** (target release expected nov 03):
  - **VDT 1.1.8-9** (Globus 2.2.4)
  - **EDG 2.0** Workload MS (RB..) and Data MS (LRC, RLS..)
  - DataTAG **GLUE Schema v 1.1**
  - **LCG-specific modifications**, adds and fixes to:
    - **LRMS ( PBS, Condor, LSF)**
    - **Globus Gatekeeper**
    - **MDS – BDII LDAP**
    - **Various Globus components**
- ◆ Reflecting into **45 different commands** in total for both end users and administrators
- ◆ Deploying a **prototypal LCG-1 tag** ( by our test's time – (beg Sept))



# LCG-CT

CertTB

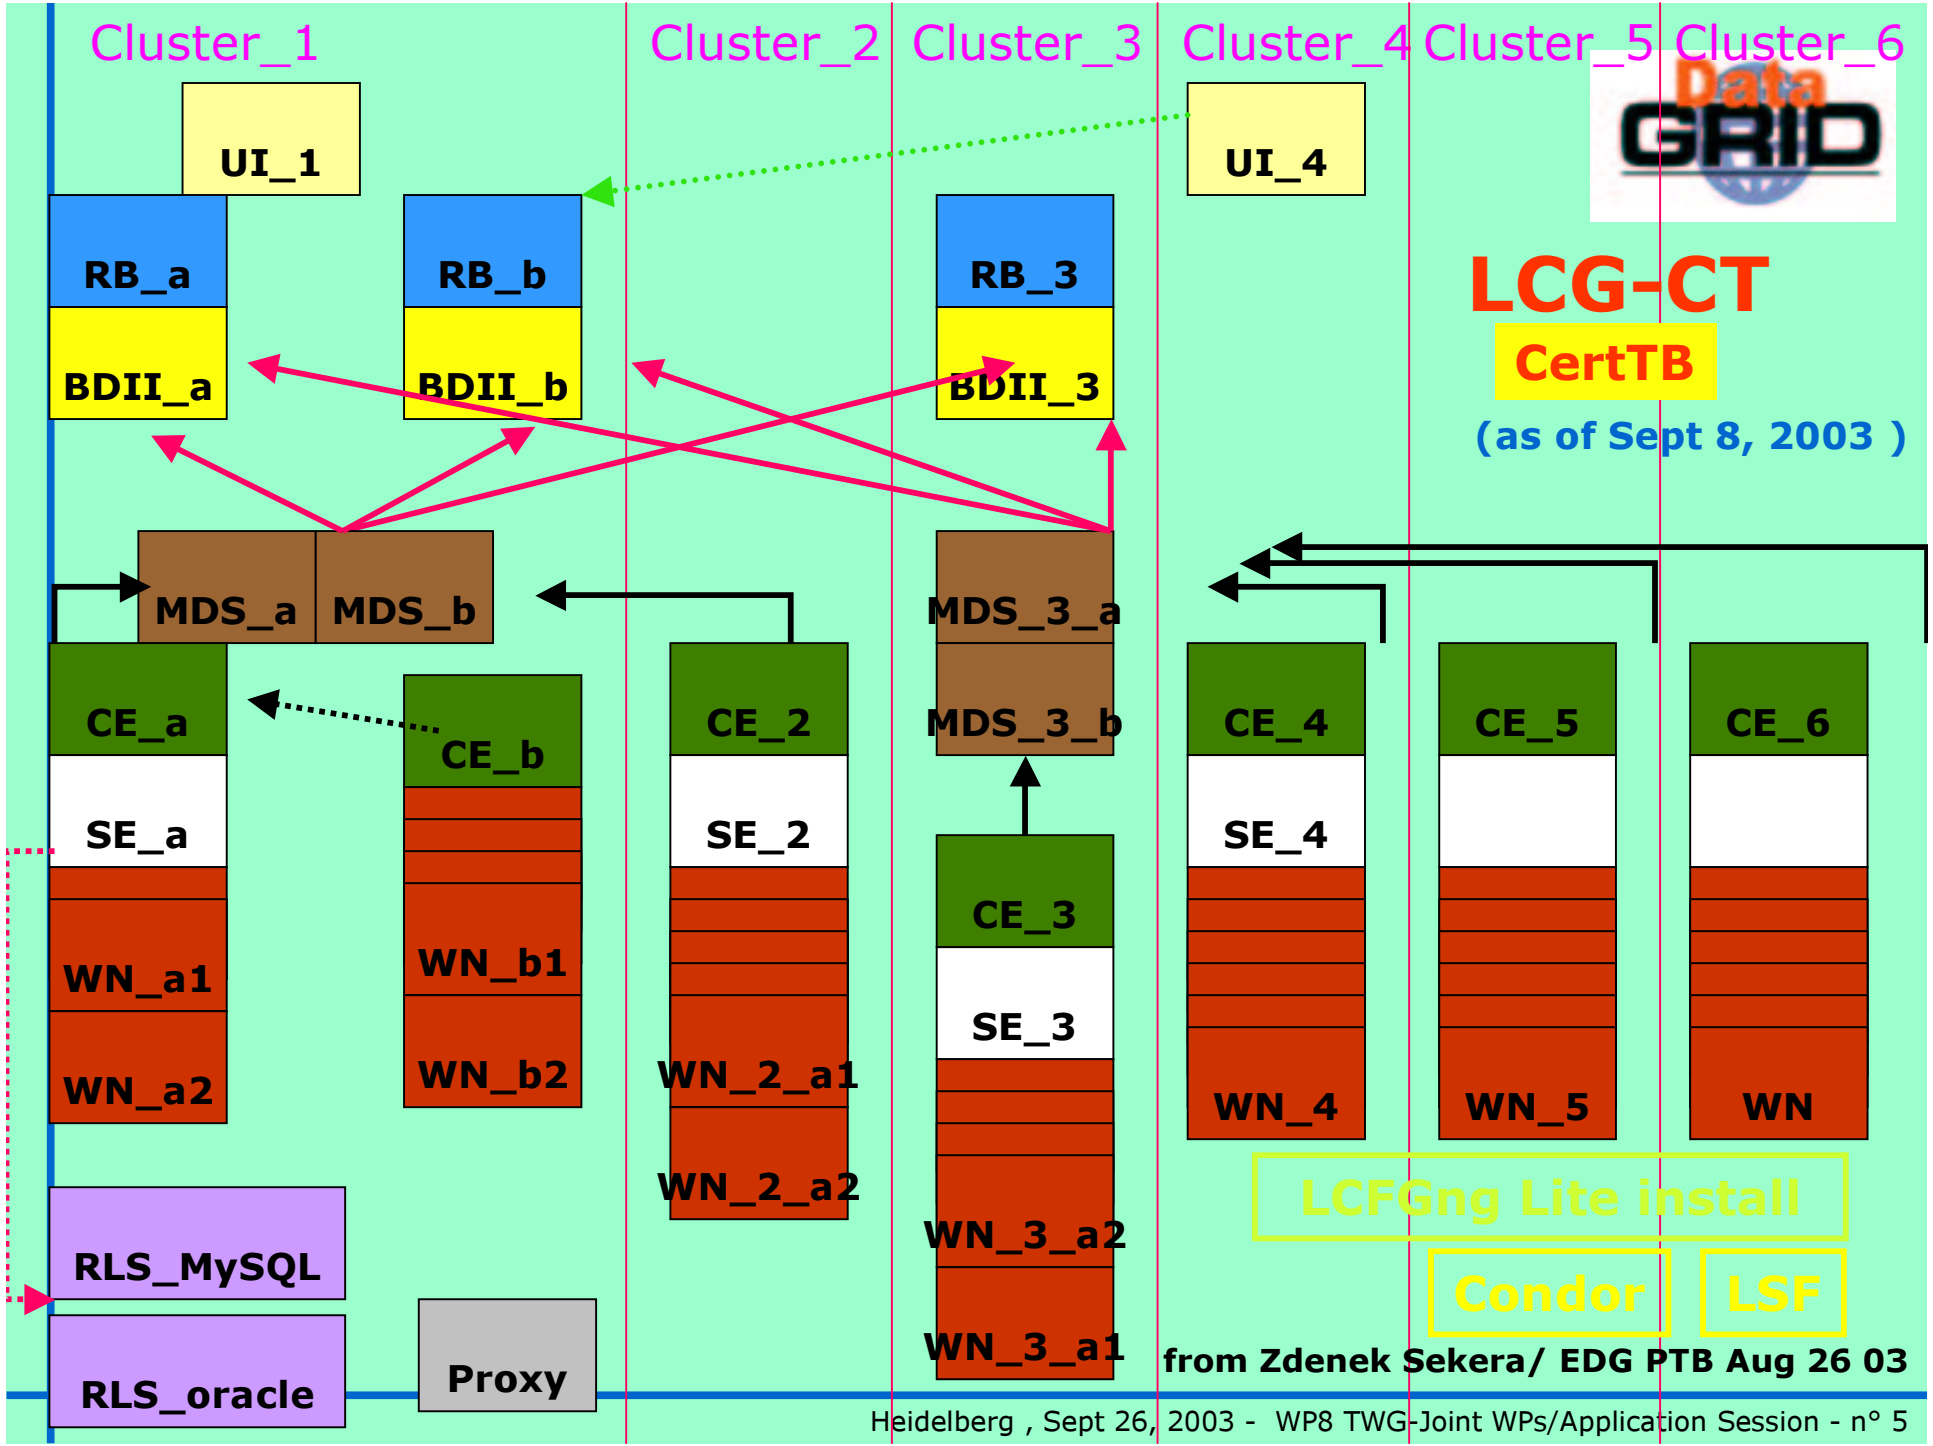
(as of Sept 8, 2003)

LCFGng Lite install

Condor

LSF

from Zdenek Sekera/ EDG PTB Aug 26 03





## Prototypal LCG-1 GRID service

- ◆ Prototypal LCG-1 **LCG GRID computing service**
- ◆ Currently **7 core sites** – (3 different continents):
  - CERN, KFKI, FZK, CNAF, FNAL, SINICA, RAL
- ◆ Truly **World-Wide** extended GRID service for LHC Computing
- ◆ Deploying **LCG1-1\_0\_0** - Rolled out on **Sept 1, 2003**.
- ◆ **RB** at CERN and CNAF, **Top-MDS** at CERN
- ◆ 4 main application VOs supported (Alice, ATLAS, CMS, LHCb)
- ◆ Opened to LCs since **Sept 15**, to experiments this week
- ◆ **Not yet** all experiment-VO s/w installed in all sites

cms:/cms 13949 1749

**GRID ICE** <http://tbed0116.cern.ch/gridice/vo/vo.php?VOname=cms>

**site: fnal.gov**

cluster: **hotdog46.fnal.gov** total phys cpus: **n/a** load5min: **n/a**

queue	free cpus	busy cpus	total cpus	run jobs	wait jobs	total jobs	max jobs
jobmanager-pbs-infinite	4	0	4	0	0	0	99999
jobmanager-pbs-long	4	0	4	0	0	0	99999
jobmanager-pbs-short	4	0	4	0	0	0	99999

storage element: **hotdog48.fnal.gov**

area	avail space (Mbytes)	used space (Mbytes)
cms:/cms	3261	1753

**site: fzk.de**

cluster: **hik-lcg-ce.fzk.de** total phys cpus: **n/a** load5min: **n/a**

queue	free cpus	busy cpus	total cpus	run jobs	wait jobs	total jobs	max jobs
jobmanager-lcgpbs-infinite	6	0	6	0	0	0	99999
jobmanager-lcgpbs-short	6	0	6	0	0	0	99999
jobmanager-lcgpbs-long	6	0	6	0	0	0	99999

storage element: **hik-lcg-se.fzk.de**

area	avail space (Mbytes)	used space (Mbytes)
cms:/cms	4618770	356

**site: grid.sinica.edu.tw**

cluster: **lcg00105.grid.sinica.edu.tw** total phys cpus: **n/a** load5min: **n/a**



## LCG-1 vs EDG-2

- ◆ DMS ROS is not foreseen in LCG
- ◆ VOMS is not yet deployed (not even in EDG-2)
- ◆ WMS-DGAS Accounting System is not deployed yet (not even in EDG-2)
- ◆ InfoSys R-GMA is not foreseen in LCG-1 ( using MDS-BDII)



edg-brokerinfo  
 edg-fetch-crl  
 edg-find-resources  
 edg-fmon-agent  
 edg-fmon-server  
 edg-gridftp-exists  
 edg-gridftp-ls  
 edg-gridftp-mkdir  
 edg-gridftp-rename  
 edg-gridftp-rm  
 edg-gridftp-rmdir  
 edg-job-attach  
 edg-job-cancel  
 edg-job-get-chkpt  
 edg-job-get-logging-info  
 edg-job-get-output  
 edg-job-list-match  
 edg-job-status  
 edg-job-submit  
 edg-local-replica-catalog  
 edg-local-replica-catalog-admin  
 edg-lrc

edg-lrc-admin  
 edg-mkgridmap  
 edg-mkgridpool  
 edg-replica-manager  
 edg-replica-manager-configure  
 edg-replica-metadata-catalog  
 edg-replica-metadata-catalog-admin  
 edg-replica-optimization  
 edg-replica-optimization-admin  
 edg-rm  
 edg-rmc  
 edg-rmc-admin  
 edg-ros  
 edg-ros-admin  
 edg-ros  
 edg-se-webservice

edg-service-mysql-setup.pl  
 edg-service-publisher-gen.pl

edg-template-conf.pl  
 edg-testbed-test

edg-vo-env

edg-voms-admin  
 edg-voms-admin-local  
 edg-voms-ldap-sync  
 edg-voms-mkgridmap  
 edg-voms-proxy-init

edg-wl-dgas-hlr-resourceInfoClient  
 edg-wl-dgas-hlr-transInfoClient  
 edg-wl-dgas-hlr-userInfoClient  
 edg-wl-grid-console-shadow

edg-wl-logev

- WMS commands
- DMS commands
- ( InfoSys commands )
- SiteAdmin and Fabric
- MassStorage-SE
- Security-VOMS

ldapsearch -h .. -p ...



# Test code sources

LCs July-Sept 03 LCG / EDG testing:



- ◆ Full general purpose generic-HEP-application test suite from Jean-Jacques Blaising: ( PERL )
  - Gets info on GRID current status (GRID config file) and creates JDL accordingly
  - Submits jobs, monitor status, retrieve output, report results
- ◆ Various LC testing JDLs and scripts (/bin/bash, PERL )
  - general and intensive stress tests
- ◆ Official EDG WP6-SC (edg-site-certification) PERL-oo test suite
  - Basic job submission, output retrieve
  - Basic Data Management tests, Filename registration and file copy
  - My Proxy tests
  - Match-Making specific tests

([http://marianne.in2p3.fr/datagrid/TestPlan/TESTSTATUS/EDG\\_20\\_TEST\\_STATUS.html](http://marianne.in2p3.fr/datagrid/TestPlan/TESTSTATUS/EDG_20_TEST_STATUS.html))

# First very preliminary results from EDG 2.0.3 **Application testbed**



- ◆ Basic Job submission, monitoring, cancel, output retrieval works ok
  
- ◆ 8 streams x 50 medium-sized jobs storm showed a small percentage ( ~3 %) of
  - Current Status: Aborted
  - Status Reason: **Cannot plan (a helper failed)** or
  
  - Current Status: Done (Cancelled)
  - Status Reason: **Cannot read JobWrapper output, both from Condor and from Maradona.**

# LCG-CT WMS tests and results (1/2)



- ◆ Tested basic job submission and output retrieve "Hello World"  
[ OK ]
- ◆ Test submission of 50 jobs with Input and Output sandboxes  
with and without retryCount=0  
[ OK ]
- ◆ Test of 250 jobs with long sleep, no resubmission (MyProxy)  
[ OK ]
- ◆ 5 job with 2GB Input&Output sandboxes, no resubmission,  
with and without parallel streams  
[ OK ]
- ◆ Match-making tests:
  - Three files on one SE, job matches to associated close CE  
[ OK ]
  - Three file on one SE, add application tag, job matches close CE  
[ OK ]
  - Access a file from a SE from job with file protocol; compute  
checksum  
[ OK ]

# LCG-CT WMS tests and results (2/2)



- ◆ Submit JJ-HEP generic application on all available GRID CEs, retrieve output, register files on iteam VO LRC via RLS from the running script on the WN, copy them on close SE.  
[ Resources previously discovered by query to InfoSys ]

[ OK ]

- ◆ Submit a numeric iterative calculation on all available GRID CEs, use BrokerInfo to find CloseSE and mount point, copy files there and register into iteam LRC through RLS

[ OK ]

# RB Stress Tests by Massive Job Submission on the LCG-CT



- ◆ RB never crashed
- ◆ ran without problems at a load of 8.0 for several days in a row 20 streams with 100 jobs each ( typical error rate  $\sim 2\%$  still present )
- ◆ RB stress test in a job storm of 50 streams , 20 jobs each :
  - 50% of the streams ran out of connections between UI and RB. (configuration parameter – but machine constraints)
  - Remaining 50% streams finished normal ( 2% error rate)
  - Time between job-submit and return of the command (acceptance by the RB) is 3.5 seconds. ( independent of number of streams )
- ◆ NOTE:

RB interrogates all suitable CE's : wide area delay-killer (interactive work) ?

# Preliminary full simulation and reconstruction tests with ALICE on the CERN LCG-CT



( beginning of Sept 03)

- ◆ Aliroot 3.09.06 (including HBT correl.) fully reconstructed events
- ◆ CPU-intensive, RAM-demanding (up to 600MB ,160MB average) ,long lasting jobs ( average 14 hours )
- ◆ Outcome:
  - > 95 % successful job submission, execution and output retrieval in a lightly loaded GRID environment
  - ~ 95 % success (**first estimate**) in a highly job-populated testbed with concurrent job submission and execution ( 2 streams of 50 AliRoot jobs and concurrent 5 streams of 200 middle-size jobs)
  - My Proxy renewal successfully exploited [ OK ]

# Currently (Sept 19) opened EDG WMS bugs ( also effecting LCG) (no show-stoppers)

<http://marianne.in2p3.fr/datagrid/bugzilla/>



- ◆ 1103 : considering working (already used) CEs in re-submission if no others available
- ◆ 1105 : mechanism should be provided to restart daemons that aren't running anymore
- ◆ 1362 : hanging LB clients if LB responsive but other RB component down – should give appropriate error message, not hang
- ◆ 1465 : purger scripts queries the LB without using a proxy (su)
- ◆ 1643 : resubmission even when proxy expired
- ◆ 1716 : in some specific cases the LB doesn't give the right job status
- ◆ 1792 : edg-job-status jobId-status max retrieval number (>1100..)
- ◆ 1798 : /etc/init.d/edg-wl-ns restart does not work (needs a sleep)
- ◆ 1770 : Jobs queues drain if the replica catalog is down



# Currently (Sept 19) opened EDG WMS bugs: 18 ( also effecting LCG ) (no show-stoppers)

<http://marianne.in2p3.fr/datagrid/bugzilla/>



- ◆ 1824: delay to get information from edgbroker-info after NS and WL daemon restart
- ◆ 1918: tmpwatch erases wm files in /tmp - /var should be used
- ◆ 1932: edg-job-get-logging-info -o --logfile prints log-info to the screen ( same for edg-job-status, list-match )
- ◆ 1934: use of commas for edg-job-\* commands working but not advertised in help menu
- ◆ 1938: Complaining lob messages from edg-wl-interlogd
- ◆ 1125: Option -c : even if -c is used, still UI seeks default config file
- ◆ ( 1843 ) : edg-find-resources doesn't have a --help option (UI, but not WP1 )
- ◆ 1848: A replacement for the old getSelectedFile method is needed given Input File users gets URL/filename (BrokerInfo)

# General issues and concerns (1/2)



- ◆ Some RGMA timeout wrongly set
  - 1 site managed to block the whole GRID for 1 day since port 8080 of MonBox wasn't accessible – It should be now fixed in RGMA
- ◆ GOUT empty quite often completely blocked all submissions
- ◆ Cryptic error messages like “A helper failed”
  - Help user to tell expectable error messages (somebody switched off a given WN) from system services critical conditions – Work on going..
- ◆ A mechanism should be provided for a sys admin to delete all jobs of a given user ( LCG-CT ) ( super user – ref bug 1465)

## General issues and concerns (2/2)



- ◆ GUI & JDL editor not fully exploited and tested yet – ( ! )
- ◆ MPI not fully exploited and tested yet – ( ! )
- ◆ Check pointing not tested yet - ( ! )
- ◆ **Gang-matching** not working at time it was tested ( 1442: **should be fixed by now** in 2.0.18, but EDG 2.0.3 includes 2.0.15-2 )
- ◆ Old relevant pending bug : wrong ClassAd ( *other.CEId* ) could block the RB – ( now **fixed** in 2.0.18 )
- ◆ Outbound connectivity for the Worker Nodes – issue for GDB



## Conclusions

- ◆ Broker seems a lot more stable, but still has some failures, and has to be tested for **scalability** : this could be the next relevant issue ( RB polling all CEs.....)
- ◆ Many new features (check pointing, interactive jobs, gang matching, output data, access cost ranking, accounting, ...) largely untested and some not integrated
- ◆ Looking forward for **VOMS integration** and **VO-based users-resources matching**