



MODIFIED RB TEST REPORT

SOFTWARE TEST AND VALIDATION REPORT

WP4 TASK4 - Verification and Quality Control

Document Filename:	CG-4.4-REP-v1.0-LIP103-modified-RB_TestReport.doc
Work package:	WP4 TASK4 - Verification and Quality Control
Partner(s):	LIP
Lead Partner:	LIP
Config ID:	CG-4.4-REP-v1.0-LIP103-modified-RB_TestReport
Document classification:	PUBLIC

Abstract: This report describes the validation performed on the package (modified) Resource Broker/Scheduling Agent developed by CrossGrid WP 3 task 3.2. The tests were performed by Mario David on the behalf of the CrossGrid task 4.4 testbed verification and quality control. The package allows the scheduling of mpich-g2 jobs.



Delivery Slip

	Name	Partner	Date	Signature
From				
Verified by				
Approved by				

Document Log

Version	Date	Summary of changes	Author
1-0-DRAFT-A	02-05-2004	Draft version	Mario David
1-0-FINAL	25-05-2004	Final version	Mario David, Jorge Gomes

CONTENTS

1. CONTEXT.....	4
1.1. TEST REQUEST	4
1.2. TEST TEAM.....	4
1.3. RESOURCES INVOLVED	4
2. TEST AND VALIDATION	5
2.1. SOFTWARE INSTALLATION	5
2.2. ADDITIONAL TESTBED MODIFICATIONS	6
2.3. TEST DEVELOPMENTS	7
2.4. USABILITY	8
2.5. FUNCTIONALITY	9
2.5.1. <i>Unit tests</i>	9
2.5.2. <i>System tests</i>	9
2.5.3. <i>Stress tests</i>	15
2.6. COMPATIBILITY	15
2.7. SECURITY AND NETWORKING.....	15
2.8. PREVIOUSLY REPORTED ISSUES	16
3. ISSUES FOUND	17
3.1. ISSUES FOUND IN THE SOFTWARE	17
3.1.1. <i>Issue 001</i>	17
3.1.2. <i>Issue 002</i>	17
3.1.3. <i>Issue 003</i>	17
3.1.4. <i>Issue 004</i>	18
3.1.5. <i>Issue 005</i>	18
3.1.6. <i>Issue 006</i>	18
3.1.7. <i>Issue 007</i>	18
3.1.8. <i>Issue 008</i>	19
3.1.9. <i>Issue 009</i>	19
3.1.10. <i>Issue 010</i>	19
3.1.11. <i>Issue 011</i>	19
3.2. ISSUES FOUND IN THE DOCUMENTATION	20
3.2.1. <i>Issue 001</i>	20
4. RECOMMENDATION	21
5. REFERENCES	22
6. INTEGRATION/VALIDATION REQUEST	23

1. CONTEXT

Test and validation of the package “modified Resource Broker” developed by CrossGrid WP 3 task 3.2. The package has a modified Scheduling Agent in order to cope with mpich-p4 and mpich-g2 parallel jobs.

1.1. TEST REQUEST

The package test request was assigned by Jorge Gomes (jorge@lip.pt) to Mario David (david@lip.pt) on 6 April 2004. The test requester was Alvaro Fernandez (alvaro.fernandez@ific.uv.es) from IFIC Valencia. The request was submitted properly through the test and validation web form, and was assigned the Request ID 10812536473.01482473305477.

The software URL provided in the “Test Request” was correct containing the several packages of the Workload Management System. The version is 2.0.18.2.1-1 is based on the WMS distributed by EDG version 2.0.18.

The URL for Installation and User Manuals were correctly provided in the Test Request. The URLs containing the Development Manual and Software Design documentation were also correctly provided.

1.2. TEST TEAM

The tests were performed by task 4.4 members from LIP Lisbon:

- Mario David (david@lip.pt)
- Jorge Gomes (jorge@lip.pt)

1.3. RESOURCES INVOLVED

The following nodes located at LIP Lisbon were involved in the tests:

Resource Broker – rb03.lip.pt

User Interface – ui01.lip.pt

Job submissions to PBS job managers at several LCG-1 sites were attempted.

IFIC	OK
CESGA	OK
CYFRONET	OK
FZK	OK
ICM	OK
UAB	OK
IFCA	FAIL
FUW	FAIL

2. TEST AND VALIDATION

2.1. SOFTWARE INSTALLATION

A small testbed comprising a LCFGng server, a RB and a UI running the LCG-1.1.4 version under the usual RH7.3 operating system, was used.

Both the RB and UI were first installed with the official LCG-1.1.4 release and job-submission tests were performed to check that the RB was running properly. Afterwards, the WMS packages corresponding to the modified RB and UI provided by the CrossGrid Task 3.2, were added to the respective rpm lists and installed on those machines.

```
UI:
cg-wp3.2-bypass-2.5.3-18
cg-wp3.2-chkpt-api-2.0.18.2.1-1
cg-wp3.2-common-api-java-2.0.18.2.1-1
cg-wp3.2-common-api-2.0.18.2.1-1
cg-wp3.2-config-2.0.18.2.1-1
cg-wp3.2-dgas-hlr-ATMClient-2.0.18.2.1-1
cg-wp3.2-dgas-hlr-jobAuthClient-2.0.18.2.1-1
cg-wp3.2-dgas-hlr-ui-2.0.18.2.1-1
cg-wp3.2-logging-api-c-2.0.18.2.1-1
cg-wp3.2-logging-api-cpp-2.0.18.2.1-1
cg-wp3.2-logging-api-sh-2.0.18.2.1-1
cg-wp3.2-services-common-2.0.18.2.1-1
cg-wp3.2-ui-api-cpp-2.0.18.2.1-1
cg-wp3.2-ui-api-java-2.0.18.2.1-1
cg-wp3.2-ui-cli-2.0.18.2.1-1

RB:
cg-wp3.2-wm-2.0.18.2.1-1.i486.rpm
cg-wp3.2-common-api-2.0.18.2.1-1.i486.rpm
cg-wp3.2-logging-api-c-2.0.18.2.1-1.i486.rpm
cg-wp3.2-logging-api-cpp-2.0.18.2.1-1.i486.rpm
cg-wp3.2-services-common-2.0.18.2.1-1.i486.rpm
cg-wp3.2-bypass-2.5.3-18.i486.rpm
cg-wp3.2-config-2.0.18.2.1-1.i486.rpm
cg-wp3.2-interactive-2.0.18.2.1-1.i486.rpm
cg-wp3.2-proxyrenewal-2.0.18.2.1-1.i486.rpm
cg-wp3.2-globus-gridftp-1.5-18.i486.rpm
cg-wp3.2-lbserver-2.0.18.2.1-1.i486.rpm
cg-wp3.2-lbserver-rgma-2.0.18.2.1-1.i486.rpm
```

In the test request the following rpm's for the UI were missing:

```
cg-wp3.2-ui-config-2.0.18.2.1-1.i486.rpm
cg-wp3.2-ui-gui-2.0.18.2.1-1.i486.rpm
```

The following rpm for the RB was missing both in the test request and in the Installation Manual [1]:

```
cg-wp3.2-locallogger-2.0.18.2.1-1
```

The CrossGrid packages were installed only after the corresponding EDG ones were removed, e.g. `edg-wl-*`. This was necessary, because the rpm's `cg-wp3.2-*` are a recompilation of the corresponding EDG ones with the source code modified to deal with `mpich-g2` jobs.

The software installs under `/opt/edg`, because the files are mostly the same as the ones distributed by EDG, except the executables corresponding to the Scheduling Agent in the RB and the `edg-job-*` executables in the UI, which had their code modified. This is an issue for the integration team, since these RPMs must replace the WMS from EDG. This implies that it is not possible to support both WMS implementations (EDG and CG) in the same user interface. This approach also violates the CrossGrid development rule that recommends that all CrossGrid developed software should be installed under `/opt/cg`, although for this particular software, an exception could be opened.

The dependencies are the same as the ones of the WMS distributed by EDG, and are all mentioned in the Installation Manual.

2.2. ADDITIONAL TESTBED MODIFICATIONS

For the RB configuration there is a manual change in the file:
`/opt/edg/etc/edg_wl.conf`

In the section corresponding to the JobController (JC), one has to add the following:

```
G2LauncherLogDir = "${EDG_WL_TMP}/jobcontrol/g2launcher";  
CheckinTimeout = 700;
```

The issue to be solved is that this file is configured through the "wlconfig" LCFGng component and eventually this component has to be modified for CG purposes. For the time being the `edg_wl.conf` will be manually modified and care has to be taken when rebooting or restarting the LCFGng components since these will overwrite this configuration file.

The Condor-G configuration file has to be modified also:
`/opt/condor/etc/condor.conf`

Namelly, one has to set:

```
ENABLE_GRID_MONITOR = FALSE
```

This can be done through LCFGng by adding into the `rb03` profile:

```
condorconfig.enable_grid_monitor    FALSE
```

2.3. TEST DEVELOPMENTS

The following JDL was used for the mpich-p4 tests:

```
JobType           = "mpich";
NodeNumber        = 2;
VirtualOrganisation = "cg";
Executable         = "test_mpi";
StdOutput         = "test_mpi.out";
StdError          = "test_mpi.err";
InputSandbox      = {"test_mpi"};
OutputSandbox     = {"test_mpi.out", "test_mpi.err"};
```

The `test_mpi` executable is a simple program to numerically calculate the value of PI. With this JDL we can select a larger number of CPU's for stress tests. One can also submit any number of such jobs.

For the mpich-g2 job submission tests, the same program was used but compiled with mpich-g2 . Initially the following JDL was used, later more nodes were added.

```
JobType           = "mpich-g2";
NodeNumber        = 2;
VirtualOrganisation = "cg";
Executable         = "cpi";
StdOutput         = "job.out";
StdError          = "job.err";
InputSandbox      = {"cpi"};
OutputSandbox     = {"job.out", "job.err"};
Rank              = other.GlueCEStateFreeCPUs;
Requirements      =
other.GlueCEUniqueID == "cg01.ific.uv.es:2119/jobmanager-pbs-long" ||
other.GlueCEUniqueID == "ce100.fzk.de:2119/jobmanager-pbs-long" ||
other.GlueCEUniqueID == "aocegrid.uab.es:2119/jobmanager-pbs-long" ;
```

Also here, a larger number of CPU's can be specified in order to test mpich-g2 between sites.

2.4. USABILITY

The software is easy to install and configure once a Resource Broker is previously up and running with the EDG/LCG middleware. This is the most difficult part.

The knowledge of how to operate a RB is essential to the understanding of the CrossGrid modified RB software.

The command line interface deployed in the UI produces the correct output with two possible exceptions which are the `edg-job-list-match` command; the output of this command is dependent on the information published by the Information Index, in particular regarding the number of free CPU's available at each site. This information has a time delay dependent on the regularity with which the Information Index is updated. This issue also affects the RB when scheduling jobs to a given site.

The command `edg-job-status`, which can be used to monitor the evolution of a job, has reduced usability with respect to the WMS deployed by EDG. The only job status produced by the Logging&Bookkeeping were:

```
Current Status:    Ready
Status Reason:    unavailable

Current Status:    Aborted
Status Reason:    Could not receive all checkins from subjobs

Current Status:    Done (Success)
Exit code:        0
Status Reason:    Job terminated successfully
```

This is probably related to the following requirement written in the Installation Manual [1], regarding the CondorG configuration and mentioned previously in section 2.2 :

```
[root@rb03 root]# less /opt/condor/etc/condor.conf
....
ENABLE_GRID_MONITOR = FALSE
....
```

The interactive response of the executables in the UI are quite acceptable, and depend more on the response from the RB and its interaction with the Information Index.

2.5. FUNCTIONALITY

The software functionality was tested through job submission to the RB, using both mpich-p4 and mpich-g2 jobs.

The RB daemons startup was successfully tested

2.5.1. Unit tests

The unit tests consisted in the successful startup of the edg-wl-* daemons in the RB. This presented no problem.

```
Starting mysqld daemon with databases from /var/lib/mysql
Starting globus-gridftp: [ OK ]
Starting JobController daemon(s)
  Starting JobController... [ OK ]
  Starting CondorG... [ OK ]
Starting LBserver: edg-wl-bkserverd
bash: /root/.bashrc: Permission denied
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
edguser  12739  0.0  0.4  4844 2276 pts/0    R   18:28   0:00 /opt/edg/sbin/edg-wl-
bkserverd -k /tmp/logging/hostproxy.pem -c /tStarting LogMonitor...
[ OK ]
Starting LocalLogger: edg-wl-interlogd and edg-wl-logd.
bash: /root/.bashrc: Permission denied
bash: /root/.bashrc: Permission denied [ OK ]
This is LocalLogger, part of Workload Management System in EU DataGrid. Copyright (c)
2002 CERN, INFN and CESNET on behalf of the EU DataGrid.
Initializing...
Parse messages for correctness...yes.
Send messages also to inter-logger...yes.
Store messages with the filename prefix "/var/tmp/dg20logd_"...yes.
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
root      12691  0.0  0.2  4260 1380 ?        S    18:28   0:00 ftpd: accepting
connections on port 2811
edguser  12770  0.0  0.8  8596
4172 ?    S    18:28   0:00 /opt/edg/bin/edg-wl-log_monitor -c edg_wl.conf
root      12776  0.0  0.2  2040 1056 pts/0    S    18:28   0:00 /bin/sh
/etc/rc.d/init.d/edg-wl-locallogger start
edguser  12790  0.0  0.3  4028 1904 pts/0    R   18:28   0:00 /opt/edg/sbin/edg-wl-
interlogd -k /tmp/logging/hostproxy.pem -c /tedguser 12794 0.0 0.2 3844 1268
pts/0    R   18:28   0:00 /opt/edg/sbin/edg-wl-logd -k /tmp/logging/hostproxy.pem -
c /tmp/loStarting NetworkServer... [ OK ]
Starting ProxyRenewal Daemon: edg-wl-renewd.
bash: /root/.bashrc: Permission denied
USER      PID %CPU %MEM    VSZ   RSS TTY      STAT START   TIME COMMAND
edguser  12834  0.0  0.3  4944 1628 pts/0    R   18:28   0:00 /opt/edg/sbin/edg-wl-
renewd -r /opt/edg/var/spool/edg-wl-renewd
starting workload manager... ok
```

2.5.2. System tests

Systems tests consists in job submission to the Resource Broker both for mpich-p4 jobs and for mpich-g2 ones. The jobs requires only 2 CPU's in either case, and are submitted to one site only.

In the first test, one mpich-p4 job submission was performed into one CE, with successful results:

```
[david@ui01 mpi]$ edg-job-submit -o jid test_mpi.jdl
Connecting to host rb03.lip.pt, port 7772
Logging to host rb03.lip.pt, port 9002
===== edg-job-submit Success =====
The job has been successfully submitted to the Network Server.
Use edg-job-status command to check job current status. Your job
identifier (edg_jobId) is:
- https://rb03.lip.pt:9000/Kaa7281NEuVetyEwdx8F5A
The edg_jobId has been saved in the following file:
/home/csyst/david/edgtest/mpi/jid
=====
```

```
*****
BOOKKEEPING INFORMATION:

Printing      status      info      for      the      Job      :
https://rb03.lip.pt:9000/Kaa7281NEuVetyEwdx8F5A
Current Status: Done (Success)
Exit code:      0
Status Reason:  Job terminated successfully
Destination:    aocegrid.uab.es:2119/jobmanager-pbs-infinite
reached on:     Tue May 18 17:38:52 2004

*****
```

The job output shown:

```
[david@ui01 mpi]$ less /tmp/jobOutput/Kaa7281NEuVetyEwdx8F5A/test_mpi.err
Process 0 on aow5grid.uab.es
Process 1 on aowlgrid.uab.es

[david@ui01 mpi]$ less /tmp/jobOutput/Kaa7281NEuVetyEwdx8F5A/test_mpi.out
pi is approximately 3.1416009869231241, Error is 0.000008333333309
wall clock time = 0.003620
```

A test requiring 8 CPUs was also successful has the following job output shows:

```
Process 0 on zeus17.cyf-kr.edu.pl
Process 5 on zeus08.cyf-kr.edu.pl
Process 4 on zeus17.cyf-kr.edu.pl
Process 1 on zeus08.cyf-kr.edu.pl
Process 2 on zeus15.cyf-kr.edu.pl
Process 3 on zeus34.cyf-kr.edu.pl
Process 6 on zeus15.cyf-kr.edu.pl
Process 7 on zeus34.cyf-kr.edu.pl
```

For the mpich-g2 job submission, care must be taken to choose only sites in which the Worker Nodes have valid IP's. The mpich-g2 job was submitted requiring 2 CPU's, and to a single site.

A successful job submission with output retrieval:

```
*****
BOOKKEEPING INFORMATION:
Printing      status      info      for      the      Job      :
https://rb03.lip.pt:9000/9YowhN0FiMBuWEC8TEHkig
Current Status:      Done (Success)
Exit code:           0
Status Reason:      Job terminated successfully
Destination:        ce.grid.cesga.es:2119/jobmanager-pbs-short
reached on:         Wed May 19 14:37:21 2004
*****
```

Job output was:

```
[david@ui01 mpi]$ less /tmp/jobOutput/9YowhN0FiMBuWEC8TEHkig/job.out0
8440: Program Started
MPI initialized
Hello from 0 of 2
0: Send_volume =      988 B/s
0: Recv_volume =       99 B/s
0: Recv_delay  =  29.867 %
...
```

However a test requiring 4 CPU's from the same site failed. In particular this site has 2 dual CPU machines and the result was the following:

```
[david@ui01 mpi]$ globus-job-run ce.grid.cesga.es /usr/bin/qstat -a
ce.grid.cesga.es:
Job ID      Username Queue      Jobname      SessID NDS TSK  Req'd Req'd  Elap
-----
932.ce.grid.ces cg017 short STDIN      31326  1 --   --   00:15 R 00:00
933.ce.grid.ces cg017 short STDIN      11422  1 --   --   00:15 R 00:00
934.ce.grid.ces cg017 short STDIN      --     1 --   --   00:15 Q  --
935.ce.grid.ces cg017 short STDIN      --     1 --   --   00:15 Q  --
[david@ui01 mpi]$ globus-job-run ce.grid.cesga.es /usr/bin/qstat -a
[david@ui01 mpi]$ edg-job-status -i jid
*****
BOOKKEEPING INFORMATION:
Printing status info for the Job : https://rb03.lip.pt:9000/gQUw-uP1FFMM-
YTs4OcGbw
Current Status:      Aborted
Status Reason:       Could not receive all checkins from subjobs
Destination:         ce.grid.cesga.es:2119/jobmanager-pbs-short
reached on:          Thu May 20 10:04:14 2004
*****
```

This means that two processes are running and the other two are in queue, but the two first are waiting for the other two to run so there is a kind of bootstrap in this case, ending in the error:

"Could not receive all checkins from subjobs".

A similar type of error has been known to exist when mpich jobs are sent to dual CPU machines. The `jobmanager-pbs` is not able to deal correctly with mpich jobs being sent to dual CPU machines. Basically it submits one job to each machine instead of using the second CPU. This issue has been discussed in the CrossGrid Iteam group and a solution is being tested at LIP.

The next test was performed requiring also 4 CPU's but running in two sites:

```
[david@ui01 mpi]$ edg-job-list-match tjdl.jdl
Connecting to host rb03.lip.pt, port 7772
*****
                        GROUPS OF CE IDs LIST
The following groups of CE(s) matching your job requirements have been
found:

      *Groups with 2 CEs*                *TotalCPUs* *FreeCPUs*
[Rank=0 TotalCPUs=11 FreeCPUs=5]
  cel100.fzk.de:2119/jobmanager-pbs-short           8           2
  cg01.ific.uv.es:2119/jobmanager-pbs-short         3           3
*****
[david@ui01 mpi]$ globus-job-run cg01.ific.uv.es /usr/bin/qstat -a
cg01.ific.uv.es:

```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	Elap S	Time
3522	cg01.ific.	cg032	short STDIN	32464	1	--	--	00:15	R	--
3523	cg01.ific.	cg032	short STDIN	20073	1	--	--	00:15	R	00:00

```

[david@ui01 mpi]$ globus-job-run cel100.fzk.de /usr/bin/qstat -a
cel100.fzk.de:

```

Job ID	Username	Queue	Jobname	SessID	NDS	TSK	Req'd Memory	Req'd Time	Elap S	Time
9061	cel100.fzk.	cg008	short STDIN	18030	1	--	--	00:10	R	--
9062	cel100.fzk.	Cg008	short STDIN	16632	1	--	--	00:10	E	00:00

In this case the job was successful:

```
Process 0 on wn200.fzk.de
Process 1 on wn100.fzk.de
Process 2 on cg06.ific.uv.es
Process 3 on cg04.ific.uv.es
```

However it was noticed that the “Destination” in the status of a job shows only the first site where the job is running, as such there is no way for the user to know where the remaining instances of the job are running. In this case we know that the job ran at both FZK and IFIC because the program produces messages to the standard error file with the number of the process instance and the node where it ran. There should be a way for a user to know where the job is actually running.

During the tests many job submissions were attempted to the PBS job managers at several LCG-1 sites. Frequently the submissions failed however in some cases they were never successful while in others sometimes they were successful sometimes they failed. The next table shows the sites where the job submission was successful at least once, and the sites where it always failed.

IFIC	SUCCESSFUL
CESGA	SUCCESSFUL
CYFRONET	SUCCESSFUL
FZK	SUCCESSFUL
ICM	SUCCESSFUL
UAB	SUCCESSFUL
IFCA	FAIL
FUW	FAIL

2.5.3. Stress tests

No stress tests were performed due to the high failure rate noticed during the system tests. These errors are likely caused by:

- PBS job-manager scheduling issues in clusters of dual CPU machines.
- Outdated CPU availability information from MDS.

2.6. COMPATIBILITY

The software is the WMS distributed by EDG and modified by CrossGrid to accept MPICH jobs, so the `edg-wl-*` rpm packages have to be replaced by the ones tested here. Apart from this issue, the software is compliant and compatible with the other components of the distribution, both EDG and CG software.

The software is not compatible with the version previously tested because the WMS version distributed by EDG (2.0.18) is not compatible with the previous tested version (1.2.21).

2.7. SECURITY AND NETWORKING

The port numbers used by the software for the (modified) Resource Broker are the same as the ones used in the EDG RB, specifically:

9000, 9001: Logging and Boocking

9002: Local logger

7772: Network Server

The package seems to conform with the CrossGrid documents in terms of security and networking.

2.8. PREVIOUSLY REPORTED ISSUES

Referring to the test report of the first version of the software [2]:

- Issue “**RPM installation**”: this issue is still present but has to be dealt inside the CG ITEAM, see section 4.
- Issue “**dg-job-status and UI information**”: has been solved in the present version.
- Issue “**dg-job-status EPERM errors**”: has been solved in the present version.
- Issue “**dg-job-list-match output for non-mpi jobs**”: has been solved in the present version.
- Issue “**JobType JDL option**”: has been solved in the present version.
- Issues “**Rank option**” and “**MDS and RB**”: this issue is still present, although it should be dealt within the CG ITEAM, see section 4.

3. ISSUES FOUND

3.1. ISSUES FOUND IN THE SOFTWARE

Some of the following issues are not directly related with the software under test but have to be dealt with inside the ITEAM, so the tester has decided to report them here since some of those issues impair the usability and the testing of the software.

3.1.1. Issue 001

(Severity: medium Priority: immediate)

The packages were installed only after the corresponding EDG ones were removed, e.g. `edg-wl-*`. This was necessary, because the rpm's `cg-wp3.2-*` are a recompilation of the corresponding EDG ones with the source code modified to accept mpich jobs. The CrossGrid ITEAM should provide complete rpm lists of the WMS middleware.

3.1.2. Issue 002

(Severity: medium Priority: immediate)

Manual configuration of the WMS had to be performed in the file:

`/opt/edg/etc/edg_wl.conf` although this file is managed by the “`wlconfig`” component of LCFGng. The changes are:

```
G2LauncherLogDir = "${EDG_WL_TMP}/jobcontrol/g2launcher";  
CheckinTimeout = 700;
```

Restarting the RB services will remove the manually added changes.

3.1.3. Issue 003

(Severity: medium Priority: immediate)

The following variable of the CondorG configuration should be added to the LCFGng profile of the RB in order to make the installation automatic, otherwise it has to be added manually.

```
condorconfig.enable_grid_monitor    FALSE
```

3.1.4. Issue 004

(Severity: medium Priority: medium)

The status provided by the `edg-job-status` has reduced information with respect to the WMS deployed by EDG. The status given by the Logging&Bookkeeping are the following:

```
Current Status:    Ready
Status Reason:    unavailable

Current Status:    Aborted
Status Reason:    Could not receive all checkins from subjobs

Current Status:    Done (Success)
Exit code:        0
Status Reason:    Job terminated successfully
```

This issue is related to the previous one.

3.1.5. Issue 005

(Severity: medium Priority: high)

The RB uses information about the number of free CPUs available at each site. It was observed that the information provided by the command `edg-job-list-match` (available CPUs) sometimes seems to be inaccurate in comparison with the information published by the GRIS CE.

3.1.6. Issue 006

(Severity: critical Priority: immediate)

The `mpich-g2` jobs only run on sites with valid IP's for the Worker Nodes, if not the job will fail. A `RunTimeEnvironment` variable at the MDS level should be used by the RB to select the capable sites for `mpich-g2` job submission. This is a similar approach to the `MPICH` flag used to flag `mpich-p4` capable sites. This is essential since a site that is capable of running `mpich-p4` may not be able to run `mpich-g2` at all. Currently several sites in the CrossGrid testbed are not capable of running `mpich-g2` jobs.

3.1.7. Issue 007

(Severity: critical Priority: immediate)

For dual CPU WN's, the PBS jobmanager is not able to deal correctly with mpich-g2 jobs. This is a known problem and is being addressed.

3.1.8. Issue 008

(Severity: medium Priority: high)

The “Destination” in the status of a job should show all the CE's in which the job is running. Preferably this should be accomplished through a new option so that by default the job status output remains compatible with the output produced by a standard broker, and only when the option is specified the complete output with the name of all CEs is shown. Optionally new commands or API functions could be provided for this purpose. This is also essential namely to detect sites that make jobs fail.

3.1.9. Issue 009

(Severity: high Priority: medium)

It is possible to submit jobs to the resource broker requesting more resources than absolute total resources available. These jobs would never be run but the resource broker accepts them. A similar extreme case is that a single job can request all the CPUs available in the testbed therefore causing a flood or denial of service. Implementing limits to the number of CPUs that a single user can request can help prevent these situations.

3.1.10. Issue 010

(Severity: high Priority: high)

During the tests was several times observed the RB reporting that “No Computing Elements” matching the job submission were found while simultaneous queries to the BDII and to the GRIS at the CEs were showing available resources. This problem should be investigated further.

3.1.11. Issue 011

(Severity: high Priority: high)

If a single process instance gets queued at one site due to insufficient resources while some other instances enter in execution reaching the MPI barrier, the whole job will fail after a timeout. This is necessary to prevent the occupation of CPUs by processes that are just waiting at the MPI barrier by instances that are in a queue somewhere (or that have been lost) and for which it is not possible to know when they will be available. However this makes essential to prevent the queuing of MPI job instances, basically a way must be found to make sure that all instances of an mpich-g2 job reach execution and are not queued.

3.2. ISSUES FOUND IN THE DOCUMENTATION

3.2.1. Issue 001

(Severity: medium Priority: high)

Packages missing in the documentation.

In the test request the following rpm's for the UI were missing:

cg-wp3.2-ui-config-2.0.18.2.1-1.i486.rpm

cg-wp3.2-ui-gui-2.0.18.2.1-1.i486.rpm

The following rpm for the RB was missing both in the test request and in the Installation Manual [1]:

cg-wp3.2-locallogger-2.0.18.2.1-1

4. RECOMMENDATION

1. The CrossGrid ITEAM should provide complete rpm lists of the WMS middleware in the next CG release.
2. The variable of the CondorG configuration should be added to the LCFGng profile of the RB:

```
condorconfig.enable_grid_monitor    FALSE
```

3. Possibility to reduce the time intervals for the BDII and GRIS information services to get more accurate information about the CPU availability in the testbed.
4. A RunTimeEnvironment variable should be set (MPICH-G2), and used by the RB to select only the sites with mpich-g2 capability. Furthermore, each site claiming to run MPICH-G2 jobs should be properly tested before entering into production.
5. The CG ITEAM has to solve the grid-job-manager-pbs problem with dual CPU machines.
6. In the edg-job-status it should be shown the complete list of CE's where the job is running, presently it only shows the first CE.

An overall recommendation is that the software can be deployed in the “Production Testbed” after implementation of the recommended actions by the CG ITEAM and developers, namely items number 1, 2 and possibly 3, 4 and 5.

5. REFERENCES

[1] Installation and User Manual:

["https://savannah.fzk.de/distribution/CrossGrid/CrossGrid/wp3/wp3_2-scheduling/docs/CG3.2-D3.3-v2.0-CSIC-installationguide.doc"](https://savannah.fzk.de/distribution/CrossGrid/CrossGrid/wp3/wp3_2-scheduling/docs/CG3.2-D3.3-v2.0-CSIC-installationguide.doc)

[2] First test report of the (modified) Resource Broker:

http://www.lip.pt/computing/projects/CrossGrid/task4/softvalidation/105811622838.5097048086212/reqest_form.html

6. INTEGRATION/VALIDATION REQUEST

The test request can be found in:

http://www.lip.pt/computing/projects/CrossGrid/task4/softvalidation/10812536473.01482473305477/request_form.html