PARALLEL RESOURCE BROKER AND USER INTERFACE TEST REPORT

SOFTWARE TEST REPORT

WP4 TASK4 - Verification and Quality Control

Document Filename:	**Testbed-middleware-testing-procedure.doc**
Work package:	**WP4 TASK4 - Verification and Quality Control**
Partner(s):	**LIP**
Lead Partner:	**LIP**
Config ID:	**CG-4.4-REP-v1.0-LIP100 ResourceBrokerTestReport**
Document classification:	**PUBLIC**

Abstract: This report describes the tests performed on the CrossGrid parallel Resource Broker and User Interface modified by task 3.2. The tests also included the CE configuration changes required by the parallel RB.

Delivery Slip

	Name	Partner	Date	Signature
From				
Verified by				
Approved by				

Document Log

Version	Date	Summary of changes	Author
1-0-DRAFT-A	05/07/2003	Draft version	M. David
1-0-DRAFT-B	09/07/2003	Draft version	J.Gomes, J.Martins, M.David
1-0-FINAL	10/07/2003	Final version	J.Gomes, J.Martins, M.David

CONTENTS

1. CONTEXT 4

1.1. TEST REQUEST 4
1.2. TEST TEAM 4
1.3. RESOURCES INVOLVED 4

2. TESTS PERFORMED 5

2.1. SOFTWARE INSTALATION 5
2.2. TESTBED MODIFICATIONS 5
2.3. SSH CONFIGURATION MODULE 6
2.4. TESTS: SUBMISSION OF MPICH-P4 JOBS 7
2.5. TESTS: SUBMISSION OF SINGLE "BATCH" JOBS AND INFORMATION SYSTEM 9
2.6. "STRESS" TEST 11
2.7. NON-MPI "STRESS" TEST 12

3. ISSUES FOUND 13

3.1. ISSUES FOUND IN THE SOFTWARE 13
3.1.1. RPM installation 13
3.1.2. Logging and bookkeeping error at submission 13
3.1.3. dg-job-status and UI information 13
3.1.4. dg-job-status "EPERM" errors 13
3.1.5. dg-job-list-match output for non-MPI jobs 14
3.1.6. JobType JDL option 14
3.1.7. Rank option 14
3.1.8. MDS and the RB 14
3.2. ISSUES FOUND IN THE DOCUMENTATION 15

4. RECOMMENDATION 16

5. REFERENCES 17

1. CONTEXT

Test of the modified parallel Resource Broker and User Interface developed within CrossGrid by WP3 task 2 "Grid resource management". The tests also included the CE configuration changes required by the parallel RB.

1.1. TEST REQUEST

A properly formatted test request form [1] was provided accompanied by an installation manual [2]. The test request was received through e-mail from Alvaro Fernandez Casani (IFIC) and authorized by the work package 4 leader Jesus Marco.

1.2. TEST TEAM

The tests were performed by task 4.4 members from LIP:

- Mário David
- João Martins
- Jorge Gomes.

1.3. RESOURCES INVOLVED

The tests involved the deployment of a dedicated test RB and UI at LIP and two clusters; one at LIP and one at IFIC. The cluster at IFIC was already configured for parallel job submission while the cluster at LIP was configured to satisfy the RB parallel job submission requirements as part of the test process. A modified UI to support the parallel RB was also deployed at LIP.

The Computing Element and the Information Index systems were "borrowed" from the "*Production Testbed"*. The next diagram shows the LIP resources used in the tests.

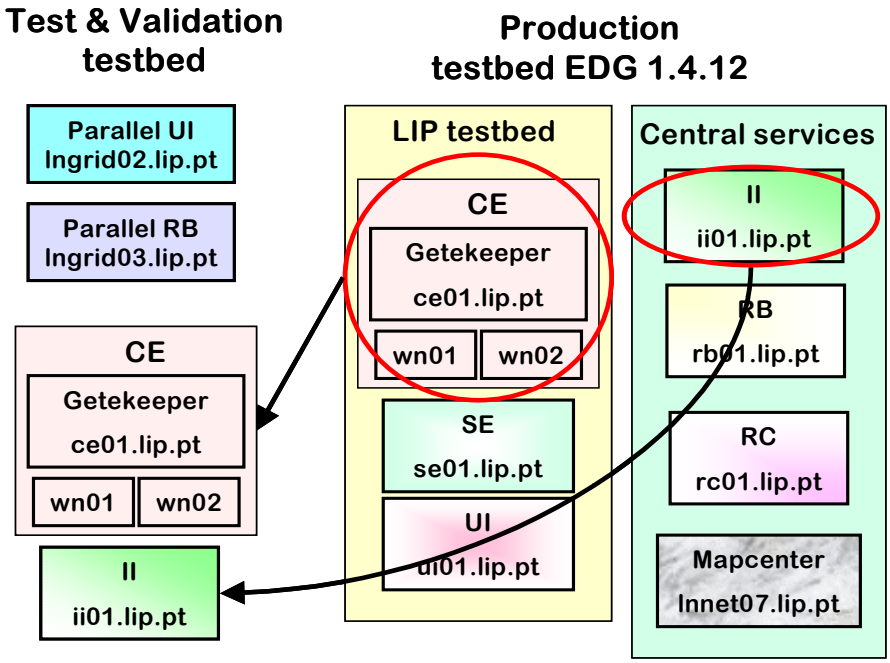

2. TESTS PERFORMED

2.1. SOFTWARE INSTALATION

Two systems were deployed at LIP, one User Interface (lngrid02.lip.pt) and one Resource Broker (lngrid03.lip.pt). Those machines were installed through LCFG from the LCFG server used to manage the Crossgrid "Prod TB" and "T&V TB". The systems were installed with RedHat 6.2 and the current middleware version EDG 1.4.12.

The following additional Crossgrid software packages were installed in the RB:

- `cg-wp3.2-jobsubmission-profile-1.2.21-1`
- `cg-wp3.2-locallogger-profile-1.2.21-1`
- `cg-wp3.2-lbserver-profile-1.2.21-1`
- `cg-wp3.2-workload-profile-1.2.21-1`
- `cg-wp3.2-logging_dev-1.2.21-1`
- `cg-wp3.2-lbserver-1.2.21-1`
- `cg-wp3.2-locallogger-1.2.21-1`
- `cg-wp3.2-jobsubmission-1.2.21-1`

While in the UI, the following were installed:

- `cg-wp3.2-userinterface-profile-1.2.21-1`
- `cg-wp3.2-workload-profile-1.2.21-1`
- `cg-wp3.2-userguide-1.2.21-1`
- `cg-wp3.2-userinterface-1.2.21-3`

The RPM's were installed with the **`--force`** option because the corresponding ones from EDG were already installed.

For these tests, an additional Worker Node was deployed in order to have a Computing Element with two single processor WN's. The LIP CE allocated for the tests was the same CE used also in the "Production TB".

The parallel RB was configured to obtain MDS information from the same Information Index used in the "Production TB".

2.2. TESTBED MODIFICATIONS

The procedures sketched in the document [2] were followed, namely:

1. *Mpich-p4 installed: exactly we need version 1.2.5. You could use another rpm at your own risk but we recommend to use the mpich-p4 rpm that can be found in crossgrid repository.*
2. *SSH special configuration: mpich needs to do ssh to other machines without being prompted for password. Ahead we'll explain two methods, but any other method you find to do this could be ok, and we'll be grateful if you send us the instructions in order to decide what is the simplest.*
3. *CE MDS additional information: our RB will only send jobs to CE's that have mpich-p4 string declared in its RunTimeEnvironment.*
4. *PBS submit script modified to allow node reservation*

2.3. SSH CONFIGURATION MODULE

Regarding the second item of section 2.2, the testers have developed a LCFG object to automatically configure the worker nodes and the computing element to allow any valid user to use SSH within the cluster without asking for a password. The LCFG module was named "sshconfig" since it configures the SSH client and servers as described in method 1 (page 5) of the document "Testbed Modifications to Support MPICH-P4 Jobs" [2]. The LCFG profile is defined as in the following example used to configure SSH for the LIP computing element, ce01.lip.pt, and worker nodes, wn01.lip.pt and wn02.lip.pt:

```
EXTRA(profile.components) sshconfig
+profile.reconfig_sshconfig start
sshconfig.hostkeys ce01 wn01 wn02
sshconfig.batchmode_ce01 yes
sshconfig.batchmode_wn01 yes
sshconfig.batchmode_wn02 yes
```

The SSH server configuration file, `sshd_config`, on each machine is thus configured for versions 1 and 2:

```
IgnoreRhosts            yes
IgnoreUserKnownHosts    yes
RSAAuthentication       yes
RhostsAuthentication    yes
RhostsRSAAuthentication yes
HostbasedAuthentication yes
```

The SSH client configuration file, `ssh_config`, on each machine is also configured for versions 1 and 2:

```
Host ce01.lip.pt wn01.lip.pt wn02.lip.pt
     BatchMode               yes
     StrictHostkeyChecking   yes
     RhostsRSAAuthentication yes
     HostbasedAuthentication yes
```

The `ssh_known_hosts` and `ssh_known_hosts2` will include, respectively the RSA1 and RSA keys of the involved hosts. Finally, the `shosts.equiv` file is filled with the hosts allowed to connect:

```
ce01.lip.pt
wn01.lip.pt
wn02.lip.pt
```

It was also noticed that, at least the worker nodes needed to allow SSH connections from the same hosts.

2.4. TESTS: SUBMISSION OF MPICH-P4 JOBS

For the job submission tests, the following JDL was used:

```
Executable     = "MPI-app-chp4";
JobType        = "mpi";
NumCPU         = 2;
StdOutput      = "std.out";
StdError       = "std.err";
InputSandbox   = {"MPI-app-chp4"};
OutputSandbox  = {"std.out","std.err"};
```

Both the executable and the JDL were obtained from the Crossgrid CVS portal at FZK:

http://gridportal.fzk.de/cgi-bin/viewcvs.cgi/crossgrid/crossgrid/wp3/wp3_2-scheduling/etc/tests/

Job submission from the LIP UI to the IFIC cluster through IFIC RB was proved successful, although some issues were found.

```
[david@lngrid02 mpitest]$ dg-job-submit -c ../uibee004.cfg -r
bee001.ific.uv.es:2119/jobmanager-pbs-qgrid -o jid mpi.jdl

Connecting to host bee004.ific.uv.es, port 7771
Logging to host bee004.ific.uv.es, port 15830
**** Warning: LB_API_LOG_EVENT ****
Unable to log Job event (ECONNREFUSED)


======================================== dg-job-submit Success
============================================
 The job has been successfully submitted to the Resource Broker.
 Use dg-job-status command to check job current status. Your job
identifier
(dg_jobId) is:

https://bee004.ific.uv.es:7846/193.136.90.175/17413370979389?bee004
.ific.uv.es:7771
 The dg_jobId has been saved in the following file:
/home/david/mpitest/jid
============================================================================
========================================
```

```
[david@lngrid02 mpitest]$ dg-job-status -c ../uibee004.cfg -i jid

Retrieving Information from LB server
https://bee004.ific.uv.es:7846
Please wait: this operation could take some seconds.

**** Error: LB_API_JOB_STATUS ****
Unable to retrieve job status for
"https://bee004.ific.uv.es:7846/193.136.90.175/104327301037049?bee0
04.ific.uv.es:7771" (ENOENT)
```

The last error is known to occur randomly in the production RB (EDG 1.4.12). The solution is to restart the **lbserver** daemon.

The job submission from the LIP UI to the IFIC cluster through the parallel LIP RB was proved successful, with the same issues referred above.

Regarding the CE at LIP, the SSH connection between the Gatekeeper and the two WN's was configured so as not to ask for the password. An **mpirun** was successfully sent from a user account in the Gatekeeper to the two WN's.

Although everything was apparently well configured for the LIP CE to be able to run mpi jobs, when jobs were submitted through the RB to the CE they failed with the following error:

```
dg-job-submit -r ce01.lip.pt:2119/jobmanager-pbs-qgrid mpi.jdl

*************************************************************
BOOKKEEPING INFORMATION:

Printing status info for the Job :
https://lngrid03.lip.pt:7846/193.136.90.175/110746302945036?lngrid03.lip
.pt:7771

Some bookkeeping information has not reached the LB server yet.
Missing information should come from UserInterface

---

dg_JobId                =
https://lngrid03.lip.pt:7846/193.136.90.175/110746302945036?lngrid03.lip
.pt:7771

Status                  =    Ready
Last Update Time (UTC)  =    Tue Jun 24 11:07:56 2003
Job Destination         =    ce01.lip.pt:2119/jobmanager-pbs-qgrid
Status Reason           =    the job failed when the job manager
attempted to run it
Job Owner               =    /C=PT/O=LIP/OU=Lisbon/CN=Mario David
Status Enter Time (UTC) =    Tue Jun 24 11:07:48 2003
Location                =    JobSubmissionService
*************************************************************
```

The only log file where some hint to the problem was found was in the CE:

```
$less /tmp/pbs_job_script.16649.err
qsub: Job exceeds queue resource limits
```

The problem was solved when the PBS server on the CE, initially configured with time shared nodes, was changed to cluster type nodes. After this change, the job submission to the cluster at LIP was proven successful.

After modifying the CE to successfully run mpi jobs sent through a RB, there was the need to verify if it's behaviour had changed when single batch jobs were sent through the "Prod TB" RB (rb01.lip.pt). Two jobs were sent to the CE through the "Prod TB" RB, but both were scheduled by the PBS server to the same WN, i.e., the PBS scheduler did not distribute the second job to the available free CPU.

The problem was traced back to the need of specifying in the PBS job submission the number of CPUs to be allocated for the job. This can be done in two ways:

- By setting the PBS server parameter **`resources_default.neednodes`** to 1. This option is not included in the initial PBS server configuration file (`pbs_server.conf`) included in the openpbs rpm package. However its usage is advised in the "Cluster installation guide" [3].
- By changing the globus PBS job submission script.

It is essential to use one of these approaches to achieve correct job scheduling and load balancing. Both approaches were tried and the following changes have been made to the globus job submission script script (**`/opt/globus/libexec/globus-script-pbs-submit`**) in the CE.
The following lines had been uncommented, see section 2.2 point 4 and page 7 of [2]:

```
# Uncomment the following lines if your pbs allows node count specification
if [ $grami_host_count -ne 0 ] ; then
  echo "#PBS -l nodes=$grami_host_count" >> $PBS_JOB_SCRIPT
fi
```

Were changed to:

```
# Uncomment the following lines if your pbs allows node count specification
if [ $grami_host_count -eq 0 ]; then
   echo "#PBS -l nodes=1" >> $PBS_JOB_SCRIPT
elif [ $grami_host_count -gt 0 ]; then
   echo "#PBS -l nodes=$grami_host_count" >> $PBS_JOB_SCRIPT
fi
```

Note: confront with the original (distributed script) and with page 7 of [2]. After this modification, two jobs sent to the CE were correctly distributed to the two WN's.

2.5. TESTS: SUBMISSION OF SINGLE "BATCH" JOBS AND INFORMATION SYSTEM

The tests of single batch jobs were undertaken with the JDL file shown below:

```
Executable     = "tsub.pl";
StdOutput      = "tsub.out";
StdError       = "tsub.err";
InputSandbox   = {"/home/david/edgtest/tsub.pl"};
OutputSandbox  = {"tsub.out","tsub.err"};
```

Note that the **JobType** option was taken out from the JDL. This executable takes about 10 minutes to run in a single processor P4 2000MHz. The job was submitted from the modified UI and scheduled by the modified RB.

A query to the Information Index running on the **ii01.lip.pt** machine was performed using the port **2170**, i.e., to the Database (BDII) containing a cache of the ldap server information. The important part of the result is shown below with respect to the LIP Gatekeeper **ce01.lip.pt**:

```
EstimatedTraversalTime: 0
FreeCPUs: 2
RunningJobs: 0
```

Although, as expected, there is one job running in the WN:

```
[root@ce01 libexec]# qstat -B
Server           Max Tot Que Run Hld Wat Trn Ext Status
---------------- --- --- --- --- --- --- --- --- ----------
ce01.lip.pt        2   1   0   1   0   0   0   0 Active
```

Performing the same query to the **2135** port of that system, i.e., to the "real" ldap server, the correct information is obtained:

```
TotalCPUs: 2
FreeCPUs: 1
NumSMPs: 0
MinSPUProcessors: 0
MaxSPUProcessors: 0
TotalJobs: 1
RunningJobs: 1
```

The *"nature"* of the BDII service as an information cache, which is updated every ten minutes from the ldap server, explains why that information is not up to date regarding a given CE.

The question now is: *"If a single non-MPI batch job is sent to the LIP CE followed by a* ***dg-job-list-match*** *for an MPI job requiring two processors, what will happen?"*

A single batch job was submitted to LIP CE and, after checking that the job was running, a **dg-job-list-match** was performed resulting in the following:

```
[david@lngrid02 mpitest]$ dg-job-list-match mpi.jdl

Connecting to host lngrid03.lip.pt, port 7771

 Groups of CEs that match job requeriments :

[Groups with 1 CEs]
   [Rank=0]  aocegrid.uab.es:2119/jobmanager-pbs-workq  TotalCPUs=3
FreeCPUs=2

   [Rank=0]  bee001.ific.uv.es:2119/jobmanager-pbs-qgrid  TotalCPUs=2
FreeCPUs=2
```

The LIP CE does not appear in the matchmaking, since it has only one free CPU, which is correct.

2.6. "STRESS" TEST

For this test the submission of multiple jobs was performed. Six MPICH-P4 jobs were submitted from the UI with the requirement for two CPU's. No resource was specified in the **dg-job-submit** command.

The result is the following: after getting the output of the six jobs, it was verified that only one did not run successfully. Two examples of the output obtained after successful completion of the jobs, are given below:

```
Output job number 1

Proceso hijo 1 ejecutado en procesador aow1grid.uab.es
Tiempo de envio de mensajes en el hijo 1 es de 0.000000 segundos
Tiempo de recepcion de mensajes en el hijo 1 es de 0.000033 segundos
Tiempo de calculo en el hijo 1 es de 0.000000 segundos
Proceso padre 0 ejecutado en procesador aow2grid.uab.es
Tiempo total de envio del padre 0.000046
Tiempo total de recepcion del padre 0.000000
Tiempo transcurrido en segundos 0.000256
Tiempo de computo del padre 0.000210
```

```
Output job number 5

Proceso hijo 1 ejecutado en procesador wn01.lip.pt
Tiempo de envio de mensajes en el hijo 1 es de 0.000000 segundos
Tiempo de recepcion de mensajes en el hijo 1 es de 0.000038 segundos
Tiempo de calculo en el hijo 1 es de 0.000000 segundos
Proceso padre 0 ejecutado en procesador wn02.lip.pt
Tiempo total de envio del padre 0.000021
Tiempo total de recepcion del padre 0.000000
Tiempo transcurrido en segundos 0.000163
Tiempo de computo del padre 0.000142
```

The output of the job, which did not run successfully, is given below:

```
Output job number 2

p0_7084:  p4_error: Timeout in making connection to remote process on
aow3grid.uab.es: 0
```

The first job was scheduled by the **broker** to run in the Univ. Barcelona cluster composed by three CPU's, therefore occupying two of them. When the second job arrived at the RB, the first one was possibly being scheduled or sent to that gatekeeper and the information published by the UAB gatekeeper was still with three free CPU's. This second job was also sent to the UAB cluster having only one free CPU, therefore explaining the error referring the third WN **aow3grid.uab.es**. This can be an issue since the last job was actually submitted even if no resources were available.

Stress tests using multiple users submitting multiple jobs were also attempted. In this case similar issues were found were multiple jobs were being scheduled to the same site staying in queue, while

resources were available at other sites. This situation is even more frequent when a ranking is specified.

2.7. NON-MPI "STRESS" TEST

Stress tests for non-MPI jobs have shown that the combination of the modified UI plus the modified RB works well for non-MPI jobs. However using the modified UI to submit jobs to a standard RB seems to be unstable since after job submission most of the times it is not possible to obtain the job status due to "EPERM" errors. For more information see section 3.1.3

3. ISSUES FOUND

3.1. ISSUES FOUND IN THE SOFTWARE

3.1.1. RPM installation

(Severity: low **Priority:** low**)**
The RPM's should not be "**`--forced`**" to install.

3.1.2. Logging and bookkeeping error at submission

(Severity: medium **Priority:** medium**)**
When submitting jobs from the modified UI an error message is obtained reporting a connection refused by the logging and bookkeeping server.

```
**** Warning: LB_API_LOG_EVENT ****
Unable to log Job event (ECONNREFUSED)
```

3.1.3. dg-job-status and UI information

(Severity: medium **Priority:** medium**)**
The **`dg-job-status`** command complains that information that should have been sent by the UI has not arrived yet. This is possibly related with the issue reported on section 3.1.2.

```
Some bookkeeping information has not reach the LB server yet.
Missing information should come from the UserInterface
```

3.1.4. dg-job-status "EPERM" errors

(Severity: high **Priority:** medium**)**

When using the modified UI to submit non-MPI jobs it is frequent to obtain an "(EPERM)" error while retrieving the job status with the **`dg-job-status`** command. Usually after retrying the **`dg-job-status`** command, the status is successfully retrieved.

However it was observed that when using the modified UI to submit non-MPI jobs to a standard EDG resource broker (rb01.lip.pt) the solution of retrying the **`dg-job-status`** command frequently does not work. Even when **`dg-job-status`** fails it is possible to remove the job with **`dg-job-cancel`**, which proves that the job was actually submitted.

This issue might be related with the logging and bookkeeping problem reported on section 3.1.2.

```
**** Error: LB_API_JOB_STATUS ****
Unable to retrieve job status for
"https://lngrid03.lip.pt:7846/193.136.90.175/18270514849110?lngrid03.lip.pt:7771" (EPERM)
```

3.1.5. dg-job-list-match output for non-MPI jobs

(Severity: low **Priority:** low**)**

When the option **`JobType = "mpi"`** is not included in the JDL, because the objective is to send a "batch" job to a single processor, the following output is obtained:

```
[david@lngrid02 edgtest]$ dg-job-list-match testsub.jdl

Connecting to host lngrid03.lip.pt, port 7771

 Groups of CEs that match job requeriments :

[Groups with 1 CEs]
   [Rank=0]     aocegrid.uab.es:2119/jobmanager-pbs-workq
TotalCPUs=65535 FreeCPUs=65535

   [Rank=0]     bee001.ific.uv.es:2119/jobmanager-pbs-qgrid
TotalCPUs=65535 FreeCPUs=65535
```

From the output it seems that thousands of free CPU's are available! Possibly the values (65535) shown indicate that the variables "TotalCPUs" and "FreeCPUs" are not relevant for a non-MPI job, however this should be CLEARLY stated in the output and in the documentation since it is misleading. Preferably for non-MPI jobs the output of **`dg-job-list-match`** should be exactly equal to the output returned by a non-modified resource broker. This can be especially important if the RB output is to be parsed by some other program.

3.1.6. JobType JDL option

(Severity: low **Priority:** low**)**

The JDL option JobType should support an option to indicate a non-MPI job submission. It should be possible to include: **`JobType = "single"`** or **`JobType = "batch"`** in the JDL to indicate a non-MPI job submission.

3.1.7. Rank option

(Severity: high **Priority:** high**)**

When the rank option is specified and multiple jobs are submitted sequentially a considerable number of jobs are scheduled to the same CE (the one with the highest rank) and stay there in queue waiting for resources while other CEs with a smaller rank are available. This might be related with a mismatch between the MDS published information and the ongoing job submissions. Possibly the jobs are scheduled to the CE while the MDS information is not yet updated. When the number of submitted jobs is high enough then jobs started to be scheduled to other CEs.

3.1.8. MDS and the RB

(Severity: high **Priority:** high**)**

The issues reported on section 2.6 and 3.1.7 should be carefully checked. In situations were multiple jobs are being submitted the information provided by the MDS might not be accurate showing

resources available that actually are in the process of being allocated. This situation may cause jobs to be submitted to clusters were there aren't enough resources.

This issue might be even more problematic with higher concurrency, when multiple Resource Brokers and users are requesting testbed resources simultaneously. In this case the information about the resource allocation may not propagate sufficiently "fast". As a consequence the RB's may have a view of the testbed that is not accurate having as a consequence job submissions to CEs that are being allocated by other jobs. In these circumstances jobs may fail or be queue even if other computing resources are free and available in the testbed.

3.2. ISSUES FOUND IN THE DOCUMENTATION

(Severity: medium **Priority:** medium**)**

The configuration of PBS in the CE as "***cluster***", must be clearly stated in the documentation as this is a requirement for proper operation of the software (even if this is usually the default).

Similarly the documentation should mention that the PBS server parameter **`resources_default.neednodes`** should be set to 1, or that the Gatekeeper PBS script (**`/opt/globus/libexec/globus-script-pbs-submit`**) should be changed. This is required for proper load balancing of non-MPI jobs.

4. RECOMMENDATION

Overall, the tests of the modified Resource Broker and User Interface software for parallel job submission were proven to work. However some logging and bookkeeping issues seem to be present and should be addressed.

Possible conflicts caused by the accuracy of the MDS published information should also be investigated. Special care must be taken upon deployment of multiple RB's sharing the whole or part of the testbed resources. The information systems queried by those RB's should be quite fast in order to avoid, as much as possible, clashing/failure of jobs being submitted to the same WN's just because the information about the resources available in a given CE was not up to date.

The main documentation issue concerns the configuration of the PBS scheduler in the Computing Element. The documentation regarding the PBS scheduler configuration should be improved.

In the current status it seems that the modified UI working with a standard RB is not stable. This situation raises the question of how many UIs will be required per site for submitting MPI and non-MPI jobs. However it seems that using the modified or the standard UI with the modified RB for running non-MPI jobs works well.

The logging of information about MPI and non-MPI jobs should be considered. It is important for measuring the testbed stability and to obtain statistics about the testbed usage, to be able to easily differentiate between MPI and non-MPI job submissions at the level of the logging and bookkeeping database.

After contacts with WP3 the task 4.4 test team was informed that the logging and bookkeeping issues might be related with the L&B version, and that other issues such as the JobType are addressed in a newer version. However this version is not yet in a format suitable for distribution (rpm) and as such could not be tested.

5. REFERENCES

[1] CROSSGRID Component Test Request Form.
Component name: **(modified) Resource Broker**, 13 June 2003

[2] “Testbed Modifications to Support MPICH-P4 Jobs”.
Document: **CG3.2-TESTBEDMODIFICATIONSMPICHP4-V1.1.DOC**

[3] "Cluster Installation Guide in CrossGRID: LCFG for Dummies", version v2.0b1,
Vangelis Floros (UOA), Cristos Markou (NCSR), Nikos Martroyiannopoulos (NCSR).
UOA - Dep. of Informatics & Telecommunications,
NCSR "Demokritos" - Institute of Nuclear Physics.