

# Table of Contents

<b>Our experiments.....</b>	<b>1</b>
Install guide for our purposes only.....	1
Patch for PBS job wrapper to leave the invocation of mpirun to the user:.....	2
Patch for LSF job wrapper to leave the invocation of mpirun to the user:.....	2
Allowing PBS to schedule all CPU's of the WN's.....	3
Setting up Lam.....	3
Random info.....	3
Known errors.....	4

# Our experiments

We shall deploy an experimental grid on these machines, the CE is called "mpitestbed":

- lxb1401.cern.ch (137.138.154.136)
- lxb1405.cern.ch (137.138.154.140)
- lxb1929.cern.ch (128.142.66.174)
- lxb1940.cern.ch (128.142.66.185)

Site-info definitions

- site-info-1b.def: default test-bed setup
- site-info-mpi-1.def: Setup for "mpitestbed" CE
- site-info-mpi-2.def: Setup for "mpicrosstest" CE

The second CE consists of the following machines and is called "mpicrosstest":

- lxb1927.cern.ch (128.142.66.172)

## Install guide for our purposes only

Disable CERN firewall:

```
rpm -e lcg-fw
/etc/init.d/iptables stop
chkconfig iptables off
```

Modify the groups file:

```
--- /etc/group.old      2006-06-21 11:12:11.000000000 +0200
+++ /etc/group         2006-06-21 11:11:47.000000000 +0200
@@ -38,7 +38,7 @@
  edguser:x:101:
  edginfo:x:102:
  rgma:x:103:
-cg:x:2688:lponcet,okeebble,maart,ycalas
+dtteam:x:2688:lponcet,okeebble,maart,ycalas
  gr:x:2648:lfield
-zp:x:1307:dqing
+atlas:x:1307:dqing
  it:x:2763:rdejong,mkoot
```

Commands to install CE+Torque+UI+WN\_torque (mostly complete...):

```
wget http://glite.web.cern.ch/glite/packages/externals/bin/rhel30/RPMS/LCG-int/glite-yaim-3.0.0-1
# OR http://grid-deployment.web.cern.ch/grid-deployment/gis/yaim/glite-yaim-3.0.0-16.noarch.rpm
rpm -Uvh glite-yaim-3.0.0-16.noarch.rpm
vi /opt/glite/yaim/functions/config_mkgridmap
/opt/glite/yaim/scripts/install_node ~rdejong/public/site-info-mpi-1.def glite-CE glite-WN glite-
ln -s /etc/init.d/grid-functions /etc/
/opt/glite/yaim/scripts/configure_node ~rdejong/public/site-info-mpi-1.def gliteCE BDII_site TORQ
```

Commands to install WN+UI:

```
wget http://glite.web.cern.ch/glite/packages/externals/bin/rhel30/RPMS/LCG-int/glite-yaim-3.0.0-1
rpm -Uvh glite-yaim-3.0.0-16.noarch.rpm
vi /opt/glite/yaim/functions/config_mkgridmap
/opt/glite/yaim/scripts/install_node ~rdejong/public/site-info-mpi-1.def glite-UI glite-WN glite-
```

```
ln -s /etc/init.d/grid-functions /etc/
/opt/glite/yaim/scripts/configure_node ~rdejong/public/site-info-mpi-1.def WN_torque UI
```

Setting up test environment with multiple MPI implementations:

- mpich2-1.0.3-1.sl3.cl.1 - Cal's RPM
- openmpi-1.0.1-1.sl3.cl.1 - Cal's RPM
- mpich-1.2.7p1-1.sl3.cl.1 - Cal's RPM
- lam-6.5.9 - Redhat Standard RPM
- lam-7.1.2-1 - currently building an RPM
- mpiexec-???

SettingUpEnvironmentVariables

## Patch for PBS job wrapper to leave the invocation of mpirun to the user:

```
--- template.mpi.pbs.sh 2006-06-21 10:35:53.000000000 +0200
+++ template.mpi.pbs.sh.new 2006-06-21 10:35:41.000000000 +0200
@@ -498,15 +498,10 @@
     fi
 fi

-HOSTFILE=${PBS_NODEFILE}
-for i in `cat $HOSTFILE`; do
-  ssh ${i} mkdir -p `pwd`
-  /usr/bin/scp -rp ./* "${i}:"`pwd`"
-  ssh ${i} chmod 755 `pwd`/${__job}
-done
-
- (
-  -cmd_line="mpirun -np ${__nodes} -machinefile ${HOSTFILE} ${__job} ${__arguments} $"
+export NRPROCS=${__nodes}
+export PBS_NODEFILE
+cmd_line="./${__job} ${__arguments} $"
+  if [ -n "${__standard_input}" ]; then
+    cmd_line="$cmd_line < ${__standard_input}"
+  fi
```

## Patch for LSF job wrapper to leave the invocation of mpirun to the user:

```
--- template.mpi.lsf.sh.2006-06-23 2006-06-23 09:18:16.000000000 +0200
+++ template.mpi.lsf.sh 2006-06-23 09:21:33.000000000 +0200
@@ -498,20 +498,10 @@
     fi
 fi

-HOSTFILE=host$$
-touch ${HOSTFILE}
-for host in ${LSB_HOSTS}
-  do echo $host >> ${HOSTFILE}
-done
-
-  -for i in `cat $HOSTFILE`; do
-    ssh ${i} mkdir -p `pwd`
-    /usr/bin/scp -rp ./* "${i}:"`pwd`"
-    ssh ${i} chmod 755 `pwd`/${__job}
-  done
-
-  (
-    -cmd_line="mpirun -np ${__nodes} -machinefile ${HOSTFILE} ${__job} ${__arguments} $"
+export NRPROCS=${__nodes}
+export LSB_HOSTS
```

```
+cmd_line="./${__job} ${__arguments} $"
if [ -n "${__standard_input}" ]; then
  cmd_line="$cmd_line < ${__standard_input}"
fi
```

## Allowing PBS to schedule all CPU's of the WN's

PBS has a config file (/var/spool/pbs/server\_priv/nodes) that lists the hosts it can use. This file is populated from the wn-list.conf file listed in site-info.def. One thing that doesn't happen, is setting the correct number of CPU's in this file. Therefore this file needs to be modified by hand. Modify the "np=1" to match the number of CPU's each WN has.

This was actually a bug in YAIM, which is now fixed.

## Setting up Lam

To test your lam setup, download the test package [☞](#). To set up the Lam environment, create a file containing the names of all the nodes to use in this run (this may be provided by Torque). Then run "recon hostfile", if you like with -d and/or -v options, to test the environment. This assumes password-less SSH logins to all machines. It dies when it receives unexpected output such as SSH banners (had to disable sending the banner in /etc/ssh/sshd\_config).

When this succeeds, run "lamboot hostfile". This performs an SSH login on all machines and starting a daemon process in the background. Make sure that port 32870 and 32871 are not firewalled on any of the machines.

After this, run the tests by untarring the test suite and running make. This builds and runs the tests using the existing lam "environment". Some errors are expected if you do not have the test suite on all machines.

BTW, don't do this as root.

## Random info

- `glite-job-submit -r lxb1405.cern.ch:2119/blah-pbs-dteam <jdl _file>`
- [EGEE-SEE mini-Grid howto](#) [☞](#)
- [Another howto](#) [☞](#)

```
# Some 1337 commands:
voms-proxy-init -voms dteam
lcg-infosites --vo dteam {ce,se,closeSE,rb,lrc,lfc}
lcg-info --list-attrs
lcg-info --list-ce --query 'Tag=*MPICH*' --attrs 'CE'
lcg-info --list-ce --query 'Tag=*MPI*' --attrs 'Tag'
glite-job-submit <jdl _file>
glite-job-status <job_ID>
glite-job-get-output <job_ID>
glite-job-list-match <jdl_file>
glite-job-cancel <job_ID>
ldapsearch -x -h lcgbdii02.gridpp.rl.ac.uk -p 2170 -b o=grid | grep -i untimeenv | grep -i mpi
```

- "gLite-3.0 lcg-RB: Condor is upgraded to 6.7.10 (...)", hence this [☞](#) should work. Maybe it's a good idea to check out the [MPI Support with Torque](#) [☞](#) at this point.
  - ◆ In Condor-6.7: "Change the mpich startup script to

```
mpirun --mca pls_rsh_agent path_to_condor_ssh
```

"

- "gLite-3.0 WLM: Support for MPI on non shared file system clusters.";
- "CE and RB never could be installed in the same computer. There are installations where RB and BDII are together."
- Add LDAP/EGEE.BDII attributes in /opt/lcg/var/gip/ldif/lcg-info-static-ce.ldif

## Known errors

- BrokerHelper: Problems during rank evaluation (e.g. GRISes down, wrong JDL rank expression, etc.)  
--> job requirements cannot be met! (see "Requirements=" or "jobType=MPICH")w

-- RichardDeJong - 14 Jun 2006

---

This topic: LCG > OurExperiments

Topic revision: r12 - 2011-06-21 - AndresAeschlimann



Copyright &© 2008-2021 by the contributing authors. All material on this collaboration platform is the property of the contributing authors.

or Ideas, requests, problems regarding TWiki? use [Discourse](#) or [Send feedback](#)