



Enabling Grids for E-science

MPI SAM Testing

John Ryan
Trinity College Dublin

MPI Workshop, CNAF
March 19th 2008

www.eu-egee.org

 The University of Dublin | Trinity College
Ollscoil Átha Cliath | Coláiste na Tríonóide




Information Society
and Media



- **A daily test was needed to verify the ability of a site to correctly execute MPI jobs**
 - To help address the usual complaint that ‘MPI is not supported’
- **Test the submission of MPI jobs (UI+RB)**
 - Use “*mpi-start* to encapsulate MPI setup...”
- **MPI-SAM test was developed (called *mpi-all* in the defs)**
 - Modified version of SAM CE tests to include a MPI test
 - MPI-SAM test consists of 3 sub-tests:
 - MPI-AD
 - MPI-START
 - MPI-INSTALL
 - All MPI flavours advertised run a simple test
 - A simple C ‘hello-world’ application
- **What MPI implementations should be supported?**
 - Those supported by *mpi-start*: OpenMPI, MPICH, MPICH2, LAM

- **MPI-AD**

- Test to see what MPI tags the site advertises
- Identifies what implementations are supported at that site
 - Want to verify that the site can do what it says


```
lcg-info --list-ce --vo $SAME_VO --query "CE=*${SAME_NODE_NAME}*" --attrs "Tag" | grep MPI
```
- Job is executing so MPICH was matched!

- **MPI-START**

- Check to verify that MPI-START is installed
 - If not installed, use bundled version + bootstrap required variables
 - *Advertise a warning in SAM results*
- Check the version of mpi-start installed
 - If (version < v0.0.52) use the bundled newer version (v0.0.58)
 - *Probably shouldn't really do this!*
 - *Advertise this fact as 'INFO' SAM alert*

- **If MPI-START not installed at the site**
 - Use the bundled version that was submitted in the input sandbox
 - Try to bootstrap for the MPI implementation installed
 - Only done for MPICH (set relevant MPI_{\$FLAVOUR}_VAR)
- **For ALL MPI flavours check:**
 1. Configuration is okay
 2. Version installed matches advertised
 3. All environment variables are set correctly
 - Any failures just show a warning in SAM results
- **Next, try to compile application (mpi-start 'pre' stage)**
 - Maybe problems with clusters that don't install compilers,
 - Info-sys solution?
- **Execute the application**
 - Num procs. requested (at submit time) = (procs. per machine +1)
- **Check the output of the application**
 - Nodes allocated is where job actually ran

- **Overall test result status**
 - OK
 - If all MPI implementations that are advertised are configured & execute correctly
 - WARN
 - If one of the advertised implementations fails
 - Site configuration is not correct
 - ERROR
 - if all MPI flavours fail to execute correctly
- **Tests are run nightly,**
 - As the DTEAM VO
 - Results available on SAM webpage

- **Nodes listed as CE in SAM DB**
 - `same-query nodename serviceabbr=CE | wc -l`
 - 484
- **Job submission (serial) possible for 331 Sites**
 - Parallel ‘MPICH’ submission for 36
- **~300 Submission failures:**
 - “Cannot plan: BrokerHelper: no compatible resources”
 - Some sites don't have enough CPUs i.e. $> (1 + \text{CPUs per machine})$
 - Not all support (or advertise) for MPI
 - `lcg-info --list-ce --vo dteam --query "Tag=MPICH*" --attrs "CE" | grep 2119 | wc -l`
 - 152
- **Of the successful submissions**
 - 14 failures (not possible to bootstrap MPICH)
 - Site may not have MPICH installed, or it's in a random location
 - 22 can run MPI jobs
 - Only 6 flawlessly, others have some config error, old version of mpi-sam etc
 - TCD fails because LAM MPI doesn't want to run

- **More languages?**
 - C++, FORTRAN
- **Separate MPI SAM tests from CE SAM**
 - So no hacking of CE SAM code + get into *lcg-sam-client-sensors*
 - Currently confusing looking at results page as MPI-SAM not critical CE test
- **MPI sensor on its own**
 - Or, maybe part of a PARALLEL applications sensor?
 - MPI variants e.g. multi-threaded MPI
- **Refine error codes currently returned**
 - What should do we mean by 'OK', 'INFO', 'WARN' 'ERROR'
 - i.e. If one of the MPI flavours fail, but others pass?