MPI Madness

Last Updated Wednesday, 19 April 2006

Did you know that there are now four popular "freely available" MPI packages for clusters.
In addition to the old standbys, MPICH and LAM/MPI, there is now MPICH2 and OPEN MPI !

Of course the question becomes what MPI is best? It depends. Most users find that one MPI
works best for different applications on their cluster but don't know which one until they try them all. Include multiple compilers (GNU, Pathscale, Intel, PGI) and now the four MPI have grown to over sixteen possible platforms. This combination can make for quite a bit of testing and headaches unless you are using our Baseline Cluster Suite (BCS). Read more
to see how we make managing MPI's, compilers, and even Sun Grid Engine a simple task.

The Way It Should Be
At Basement Supercomputing we understand clusters because we use clusters. We know that
trying to manage multiple versions of MPI is a real pain. And, once you get your environment setup to compile and link your code with an MPI, then you have to
make sure it works with your scheduler. Our solution is to use best in class
open software solutions to solve these and other cluster problems.

As an example, with our Baseline Cluster Suite, uses can change MPI environments by simply
typing the following:

```
module load MPI/mpich-gnu4
```

That is it, use mpicc, mpif77, or mpirun and it just works. Want to submit a job to Grid Engine, simple, a tested template script is ready to use. Suppose you want to try another environment, just type the following:

```
module rm mpi/mpich-gnu4
module load pgi
module load mpi/lam-pgi
```

You are now using LAM/MPI and the Portland Group compilers. And, the Grid Engine script is ready and waiting.

The key to this flexibility is the integration of the Modules package into our tool-set. We also have
prepared Grid Engine scripts for each MPI. The important thing to remember is
Basement Supercomputing has integrated and tested these tools so you don't have to waste time figuring out how to do it yourself. Finally, if you have a problem, support is
just a click away.

You do the science, We do the cluster

Contact
us today!