

Introduction to Cheetah

Rebecca Fahey
faheyra@ornl.gov

The submitted manuscript has been authored by a contractor of the U.S. Government under Contract No. DE-AC05-00OR22728. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes.

Contents

- Hardware
- Connecting
- Shell problems
- Compilers
- Interactive use
- LoadLeveler

Hardware Overview

- 27, IBM Power4 nodes
- 32, 1.3 GHz, Power4 processors per physical node (864 processors)
- Over 4.5 Teraflops in the compute partition
- 3 nodes reserved (64 processors)



Connecting

- Method:
 - Use SSH from a Linux or UNIX workstation and Putty, Cygwin, or other such program, from a PC.
- Commands:
 - `ssh user@cheetah.ccs.ornl.gov`
 - Options you may need include:
 - x, -1, -2, -v

Secure X-Windows

- If using Cygwin on one of the PCs here:
 - Before sshing to Cheetah, issue the “startx” command
 - From one of the X-windows, issue the following ssh command:
`ssh -X user@cheetah.ccs.ornl.gov`

Default shell

- Current shell is csh
- To change to ksh

```
>dcecp
dcecp> account modify user -shell /usr/bin/ksh
```
- To change to tcsh

```
>dcecp
dcecp> account modify user -shell /usr/local/bin/tcsh
```

Backspace problems

- If your backspace key fails to work, try:
`stty erase <press backspace>`
- In `.cshrc` or `.profile`, try one of the following:
`stty erase ^? (literally)`
or `stty erase ^H`

File Systems

- Home directories:
 - Located at `/dfs/home/<username>`
 - 100 MB quota
 - Subdirectories: `bin`, `private` `public` `www` `yesterday`
- Work directories
 - Located at `/tmp/gpfs750a/<username>`
 - No quota limit
 - Use `$$SYSTEM_USERDIR` to access

Compilers

- IBM Fortran
 - `xlf`, `xlf90`, `xlf95`
 - `mpxf`, `mpxf90`, `mpxf95` (for MPI codes)
 - Add `_r` suffix to any of the above for thread safe libraries. Use for MPI-I/O, 64-bit, LAPI and threaded codes.
- IBM C/C++
 - C: `xlc`, `cc` (more lenient, error->warning)
 - C++: `xlc`, `CC`
 - Add `mp` prefix for MPI codes and `_r` suffix for thread safe libraries
- GNU
 - `gcc`, `g++`

Interactive Use

- Start an interactive job with the “poe” command. Possible options to this command include:

“poe” option	Description
<code>-procs n</code>	The number (“ <i>n</i> ”) of parallel processors. Use with either “ <code>-tasks_per_node</code> ” or “ <code>-nodes</code> ”
<code>-nodes n</code>	The number (“ <i>n</i> ”) of nodes. Use with either “ <code>-procs</code> ” or “ <code>-tasks_per_node</code> .”
<code>-tasks_per_node n</code>	The number (“ <i>n</i> ”) of parallel processes per node. Use with either “ <code>-procs</code> ” or “ <code>-nodes</code> .”
<code>-rmpool n</code> (required)	The resource-manager pool that LoadLeveler will use to allocate nodes. The nodes of the compute partition are in pool “1.”

Interactive Use Cont.

- Set the environment variable `LOADL_INTERACTIVE_CLASS` to “workshop”
 - `setenv LOADL_INTERACTIVE_CLASS workshop`
 - `export LOADL_INTERACTIVE_CLASS=workshop`
- To start an interactive job
 - `poe a.out -rmpool 1 -procs 4 -nodes 1`
This runs `a.out` on 4 processors on the same node
- If the processors are not available, this will generate a failure
- For a large number of processors or when the system is busy submit a batch job

MPI batch file

```
#!/bin/ksh
#@ class = workshop
#@ job_type = parallel
#@ network.MPI = ccss,shared,US
  or #@ network.LAPI = ccss,shared,US
#@ output = $(host).$(jobid).out
#@ error = $(host).$(jobid).err
#@ wall_clock_limit = 30:00
#@ tasks_per_node = 4
#@ node = 1
#@ environment = var1=<setting>; var2=<setting>
#@ queue
pwd
export MP_SHARED_MEMORY=yes
poe a.out
```

OpenMP batch file

```
#!/bin/ksh
#@ class = workshop
#@ job_type = serial
#@ output = $(host).$(jobid).out
#@ error = $(host).$(jobid).err
#@ wall_clock_limit = 30:00
#@ resources = ConsumableCpus(8)
#@ queue
pwd
export OMP_NUM_THREADS=8
a.out
```

Specifying Memory

To request more than 1 GB of memory per processor add the following to your batch script:

```
- #@ resources = ConsumableMemory(2 gb)
or
- #@ resource = ConsumableCpus(4)
ConsumableMemory(2 gb)
```

Controlling jobs

- Submit a job: `llsubmit script`
- Stopping a job: `llcancel jobname`
 - Deletes queued jobs
 - Halts running jobs

Job status: llq

```
cheetah0033% llq
Id                               Owner      Submitted   ST PRI Class      Running On
-----
cheetah0033.7871.0               ernie      7/23 09:24 R 50 royalty cheetah0049
cheetah0033.7869.0               snuffy    7/23 08:44 R 50 batch   cheetah1060
cheetah0033.7874.0               zoe       7/23 09:52 R 50 batch   cheetah0097
cheetah0033.7910.0               bob       7/23 21:02 R 50 batch   cheetah0273
cheetah0033.7911.0               gordon    7/23 21:06 R 50 batch   cheetah1122
cheetah0033.7886.0               oscar     7/23 13:44 I 50 No_Class
cheetah0033.7896.0               bob       7/23 19:38 I 50 batch
```

Questions?

- <http://www.ccs.ornl.gov/Cheetah/filesystems.html>
- <http://www.ccs.ornl.gov/Cheetah/interactive.html>
- <http://www.ccs.ornl.gov/Cheetah/LL.html>
- `consult@ccs.ornl.gov`