Using UNCC Parallel Programming Cluster

B. Wilkinson: Modification date January 23, 2014

Overview

The current state of the cluster is shown below:

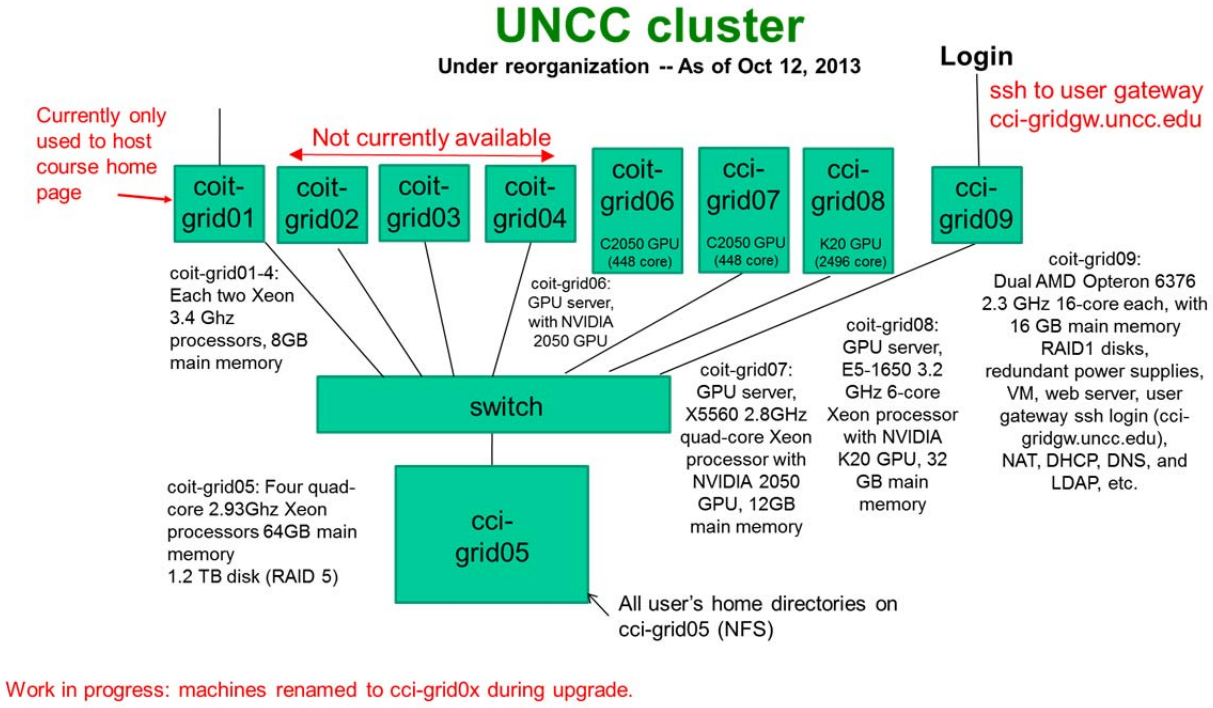

Note this cluster does not use a job scheduler. Your programs can be executed directly for the command line and you will have direct interactive access to your programs. However, you will not get control of the command line until the program finishes (or you issue a control-C). If necessary, you can run your jobs in the background with the & option on the command line.

FOR THE SAKE OF OTHERS, please make sure you do not leave any long running processes.

VERY VERY IMPORTANT: Do not ssh *from* internal nodes 05, 07 or 08 to anywhere, only from cci-gridgw (09) as currently ssh *from* internal nodes 05, 07 or 08 causes really significant problems including loosing all your files. ONLY run MPI from **cci-gridgw.uncc.edu**. If you need to use only an internal node, specify the node in a machines file.

Backup — There is no automatic file backup on this cluster although a RAID-5 system is in place for home directories. *It is highly recommended that keep a back-up copy of your files on another computer.*

Logging on

Users ssh to **cci-gridgw.uncc.edu** (user gateway on **cci-grid09**). When you log in for the first time, accept all default settings (no key passphase, password-less ssh between nodes):

```
login as: xxx
xxx@cci-gridgw.uncc.edu's password:
Creating directory '/nfs-home/xxx'.

We need to set up some things so you can SSH between
nodes securely from your account.

Please hit enter a bunch, accepting all defaults!

Generating public/private rsa key pair.
Enter file in which to save the key (/nfs-home/xxx/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /nfs-home/xxx/.ssh/id_rsa.
Your public key has been saved in /nfs-home/xxx/.ssh/id_rsa.pub.
The key fingerprint is: …

All done!

You should now be able to run the following commands
to connect to the other three nodes:

ssh cci-grid05
ssh cci-grid07
ssh cci-grid08

Enjoy!

[xxx@cci-gridgw ~]$
```

Home directory — The full path is **/nsf-home/<username>/...** .

Executing programs that exist in current directory — As is usual on Linux systems, the current directory is not in the path and it is necessary to provide **./** or the full path otherwise, you will get a "File not found" error.

X11 libraries: Currently only installed on **cci-gridgw.uncc.edu** and **cci-grid05**. Compile with **–lX11**. See separate notes on creating graphical output for more details.

MPI

MPICH2 – Compute nodes installed on 05, 07, 08. Run MPI from **cci-gridgw.uncc.edu** only.

Compile **hello.c** with **mpicc**.

mpicc –o hello hello.c

The UNCC cci-grid0x cluster uses the MPI Scalable Process Management System (Hydra). To execute MPI programs, use the **mpiexec.hydra** command (not **mpiexec**, which does not work). Execute **hello** with the Hydra process manager with:

mpiexec.hydra -n <*number of processes*> ./hello

Without a machines file, the program will run just on **cci-gridgw.uncc.edu** (grid09)

Machines file:

Compute nodes communicate through internal names. Currently only compute nodes **cci-grid05**, **cci-grid07** and **cci-grid08** are available. A machines file to use these nodes and the front node of the cci-grid0x cluster would be:

```
cci-grid05
cci-grid07
cci-grid08
cci-gridgw.uncc.edu
```

Then:

```
mpiexec.hydra -machinefile machines -n 4 ./prog
```

would run **prog** with four processes, one on **cci-grid05**, one on **cci-grid07**, one on **cci-grid08**, and one on **cci-gridgw.uncc.edu**

The machines file can include how many processes to execute on each computer, for example:

```
                        # a comment
cci-grid05:2            # first 2 processes on 05
cci-grid07:3            # next 3 processes on 07
cci-grid08:4            # next 4 processes on 08
cci-gridgw.uncc.edu:1   # Last process on gridgw (09)
```

(10 processes in total). To execute as specified, issue:

```
mpiexec.hydra -machinefile machines -n 10 ./prog
```

If more than ten processes were specified, they would be scheduled in round robin fashion.

More information on Hydra Commands:
http://software.intel.com/sites/products/documentation/hpc/ics/impi/41/lin/Reference_Manual/index.htm#Global_Options.htm Some commands do not seem to work (-hosts, -machine)

OpenMP

You are allowed to ssh to internal nodes to execute OpenMP.

Compile an OpenMP program **omp_hello.c** with the regular gcc compiler (version 4.2 onwards) using the command:

```
cc -fopenmp -o omp_hello omp_hello.c
```

Execute as a regular executable:

./omp_hello.c

Combined (hybrid) OpenMP/MPI programs

Compile hybrid program **omp-mpi_hello.c** with:

mpicc -fopenmp -o omp-mpi_*hello* omp-mpi_hello.c

Note this command invokes the regular cc compiler.

Execute as an MPI program on **cci-gridgw.uncc.edu** using:

mpiexec.hydra -machinefile <*machinesfile*> -n <*number of processes*> ./omp-mpi_*hello*