

EPL372

Hands on training session

Thekla Loizou
The Cyprus Institute

Overview

- Cy-Tera
- Euclid
- How to access Euclid (ssh)
- Software Environment (Modules)
- Hands on exercises



Cy-Tera

- Hybrid CPU/GPU Linux Cluster
- Computational Power
 - 98 x 2 x Westmere 6-core compute nodes
 - 18 x 2 x Westmere 6-core + 2 x NVIDIA M2070 GPU nodes
- Theoretical Peak Performance (TPP) = 30.5Tflops
 - Each compute node = 128GFlops
 - Each GPU node = 1 Tflop
- 48 GB memory per node
- 40Gbps QDR Infiniband
- Storage: 360TB raw disk



Euclid

- Training Cluster
- Hybrid CPU/GPU Linux Cluster
- Computational Power
 - 6 eight-core compute nodes + 2 NVIDIA Tesla T10 processors
- Theoretical Peak Performance (TPP) = 0.5 Tflop/s
- 16 GB memory per node
- Infiniband Network

How to access Euclid

- To access the system you must have access to your private key:
 - If the key is stored in your .ssh folder in your home directory:

```
ssh username@euclid.cyi.ac.cy
```

- If the key is stored in another directory or in a usb:

```
ssh -i /key_path/id_rsa username@euclid.cyi.ac.cy
```

- Your key should only be readable and writable by you. If not, please go to the directory where your key is stored and change its permissions accordingly:

```
chmod 600 id_rsa
```

Environment Setup with Modules

- The software environment used on Cyl systems can be managed via Modules
- Modules facilitate the task of updating applications and provide a user-controllable mechanism for accessing software revisions and controlling combination of versions
- Useful module commands
 - `module avail` : list all available modules
 - `module load X` : load module X
 - `module unload X` : unload module X
 - `module list` : list currently loaded modules
 - `module whatis X` : display information for module X
 - `module purge` : unload all loaded modules

Modules on Euclid

----- /fhgfs/euclid/buildsets/eb141014/modules/all -----	
ATLAS/3.10.1-gompi-1.5.12-LAPACK-3.4.2	
Autoconf/2.69-GCC-4.8.2	
Autoconf/2.69-gccuda-2.6.10	(D)
Automake/1.14-GCC-4.8.2	
Automake/1.14-gccuda-2.6.10	(D)
Bison/2.5-GCC-4.6.3	
Bison/2.5-goolf-1.4.10	
Bison/2.7-GCC-4.7.3	
Bison/2.7-goolf-1.4.10	
Bison/2.7-goolf-1.6.10	
Bison/2.7-ictce-5.3.0	(D)
CMake/2.8.4-goolf-1.4.10	
CMake/2.8.12-goolf-1.4.10	
CMake/2.8.12-goolfc-2.6.10	(D)
CUDA/5.0.35-1	
CUDA/5.5.22-GCC-4.8.2	
CUDA/5.5.22	(D)
cosmoLoPy/0.1.104-goolf-1.4.10-Python-2.7.3	
Cube/4.2-goolf-1.4.10	
Cube/4.2.3-gompi-1.6.10	
Cube/4.2.3-ictce-5.3.0	(D)
Doxygen/1.8.1.1-goolf-1.4.10	
Doxygen/1.8.3.1-goolf-1.4.10	
Doxygen/1.8.3.1-ictce-5.3.0	(D)
EasyBuild/1.10.0	
EasyBuild/1.15.2	(D)
FFTW/3.3.3-gmvapich2-1.7.12	
FFTW/3.3.3-gompi-1.3.12	
FFTW/3.3.3-gompi-1.4.10	
FFTW/3.3.3-gompi-1.6.10	
FFTW/3.3.3-gompi-2.6.10	(D)
GCC/4.6.3	
GCC/4.6.4	
GCC/4.7.2	
GCC/4.7.3	
GCC/4.8.1	
GCC/4.8.2	
GCC/4.8.3	
GCC/4.9.2	(D)
GDL/0.9.3-goolf-1.4.10	
GLib/2.34.3-gompi-1.6.10	
GLib/2.34.3-goolf-1.4.10	
GLib/2.34.3-ictce-5.3.0	(D)
GROMACS/4.6.1-goolf-1.4.10-hybrid	
GROMACS/4.6.5-goolf-1.4.10-hybrid	
GROMACS/4.6.5-goolf-1.4.10-mt	(D)
GSL/1.16-goolf-1.4.10	
HDF5/1.8.9-goolf-1.4.10	
HDF5/1.8.10-patch1-goolf-1.4.10	
HDF5/1.8.10-ictce-5.3.0-gpfs	(D)
HPL/2.0-goolf-1.4.10	
HPL/2.0-ictce-5.3.0	(D)
IOR/2.10.3-goolf-1.4.10-mpiio	
JasPer/1.900.1-goolf-1.4.10	
OTF2/1.4-ictce-5.3.0	
OpenBLAS/0.2.6-gompi-1.3.12-LAPACK-3.4.2	
OpenBLAS/0.2.6-gompi-1.4.10-LAPACK-3.4.2	
OpenBLAS/0.2.8-gompi-1.5.12-LAPACK-3.4.2	
OpenBLAS/0.2.8-gompi-1.6.10-LAPACK-3.4.2	
OpenBLAS/0.2.8-gompi-2.6.10-LAPACK-3.4.2	(D)
OpenMPI/1.6.4-GCC-4.6.4	
OpenMPI/1.6.4-GCC-4.7.2	
OpenMPI/1.6.5-GCC-4.8.1	
OpenMPI/1.7.3-GCC-4.8.2	
OpenMPI/1.7.3-gccuda-2.6.10	(D)
OpenMPI/1.8.1-GCC-4.8.3	
OpenMPI/1.8.4-GCC-4.9.2	(D)
PAPI/5.2.0-ictce-5.3.0	
PAPI/5.3.2-gompi-1.6.10	
PAPI/5.3.2-ictce-5.3.0	(D)
PDT/3.19-gompi-1.6.10	
PDT/3.19-ictce-5.3.0	(D)
PP/1.6.4-goolf-1.4.10-Python-2.7.3	
PP/1.6.4-ictce-5.3.0-Python-2.7.3	(D)
Python/2.5.6-gompi-1.6.10-bare	
Python/2.5.6-goolf-1.4.10-bare	
Python/2.7.3-goolf-1.4.10	(D)
Python/2.7.3-ictce-4.1.13	
Python/2.7.3-ictce-5.3.0	(D)
Python/3.2.3-goolf-1.6.10	
Python/3.3.2-goolf-1.4.10	
QT/4.8.4-goolf-1.4.10	
QT/4.8.5-gompi-1.6.10	
QT/4.8.5-ictce-5.3.0	(D)
SIONlib/1.5.2-gompi-1.6.10	
SIONlib/1.5.2-gompi-2015a	
ScalaLAPACK/2.0.2-gompi-1.3.12-OpenBLAS-0.2.6-LAPACK-3.4.2	
ScalaLAPACK/2.0.2-gompi-1.4.10-OpenBLAS-0.2.6-LAPACK-3.4.2	
ScalaLAPACK/2.0.2-gompi-1.6.10-OpenBLAS-0.2.8-LAPACK-3.4.2	
ScalaLAPACK/2.0.2-gompi-2.6.10-OpenBLAS-0.2.8-LAPACK-3.4.2	(D)
ScalaSca/2.0-goolf-1.4.10	
ScalaSca/2.1-ictce-5.3.0	(D)
Score-P/1.3-ictce-5.3.0	(D)
Szip/2.1-goolf-1.4.10	
Szip/2.1-ictce-5.3.0	(D)
Tcl/8.5.12	
Valgrind/3.8.1-goolf-1.4.10	(D)
WRF/3.4-goolf-1.4.10-dmpar	
WRF/3.4-ictce-5.3.0-dmpar	(D)
arpack-ng/3.1.3-ictce-5.3.0	
astropy/1.0.6-goolf-1.4.10-Python-2.7.3	(D)
astropy/1.0.6-goolf-1.4.10-Python-3.3.2	
binutils/2.24	
bzip2/1.0.6-gompi-1.6.10	
bzip2/1.0.6-goolf-1.4.10	
bzip2/1.0.6-goolf-1.6.10	(D)
bzip2/1.0.6-ictce-4.1.13	
bzip2/1.0.6-ictce-5.3.0	(D)
hello/0.1-ictce-5.3.0	(D)
hwloc/1.6.2-GCC-4.6.4	
hwloc/1.6.2-GCC-4.7.2	
hwloc/1.7.2-GCC-4.8.2	
hwloc/1.8-gccuda-2.6.10	
hwloc/1.9-GCC-4.8.3	(D)
hwloc/1.10.0-GCC-4.9.2	(D)
icc/2011.13.367	
icc/2013.3.163	(D)
icctfort/2011.13.367	
icctfort/2013.3.163	(D)
ictce/4.1.13	
ictce/5.3.0	(D)
ifort/2011.13.367	
ifort/2013.3.163	(D)
impi/4.1.13	
impi/5.3.0	(D)
imkl/10.3.12.361-impi-4.1.13	
imkl/11.0.3.163-impi-5.3.0	(D)
impi/4.1.0.027-icctfort-2011.13.367	
impi/4.1.0.030-icctfort-2013.3.163	(D)
impi/4.1.0.030	
inputproto/2.3-ictce-5.3.0	
ipp/7.0.5.233	
itbb/4.2.5.176	
kbproto/1.0.6-ictce-5.3.0	
libX11/1.6.1-ictce-5.3.0	
libXau/1.0.8-ictce-5.3.0	
libffi/3.0.13-gompi-1.6.10	
libffi/3.0.13-goolf-1.4.10	
libffi/3.0.13-ictce-5.3.0	(D)
libpng/1.5.13-goolf-1.4.10	
libpng/1.5.13-ictce-5.3.0	(D)
libpng/1.6.0-goolf-1.4.10	
libreadline/6.2-gompi-1.6.10	
libreadline/6.2-goolf-1.4.10	
libreadline/6.2-goolf-1.6.10	
libreadline/6.2-ictce-4.1.13	
libreadline/6.2-ictce-5.3.0	(D)
libxcb/1.8-ictce-5.3.0-Python-2.7.3	
make/3.82-goolf-1.4.10	
matplotlib/1.2.1-goolf-1.4.10-Python-2.7.3	(D)
matplotlib/1.2.1-ictce-5.3.0-Python-2.7.3	
matplotlib/1.3.1-goolf-1.4.10-Python-3.3.2	(D)
mpil4py/1.3-goolf-1.4.10-Python-2.7.3	
ncurses/5.9-gompi-1.6.10	
ncurses/5.9-goolf-1.4.10	
ncurses/5.9-goolf-1.6.10	
ncurses/5.9-goolfc-2.6.10	
ncurses/5.9-ictce-4.1.13	
ncurses/5.9-ictce-5.3.0	(D)
netCDF/4.2.1.1-goolf-1.4.10	
netCDF/4.2.1.1-ictce-5.3.0	(D)

Hands on exercises

- Copy the hands on exercises directory into your home directory:
 - If you are not in your home directory first type “cd”

```
cp -r /opt/exercises .
```

- Go into the exercises directory:

```
cd exercises
```


Steps to run a code

1. Load the modules needed to compile the code
2. Compile the code
3. Create a job script
 - The job script will be submitted to the queue
 - Job script asks for resources to run the compiled code
 - Through the job script the code will run on the compute nodes
4. Submit job script

Exercise 1- MPI Hello World

```
#include <stdio.h>
#include <mpi.h>
int main(int argc, char ** argv) {
    int size,rank;
    int length;
    char name[80];

    MPI_Init(&argc, &argv);
    MPI_Comm_rank(MPI_COMM_WORLD,&rank);
    MPI_Comm_size(MPI_COMM_WORLD,&size);
    MPI_Get_processor_name(name,&length);

    printf("Hello MPI World! Process %d out of %d on %s\n",rank,size,name);
    MPI_Finalize();
    return 0;
}
```

Exercise 1 - Compile the code

- Load the goolf module (g - GCC compiler, o - OpenMPI, o - OpenBLAS, l - ScaLAPACK, f - FFTW)

```
module load goolf
```

- Compile the code

```
mpicc hello.c -o hello
```

Exercise 1 - Create job script

```
#!/bin/bash

#SBATCH --job-name=hello-world
#SBATCH --nodes=2
#SBATCH --ntasks-per-node=8
#SBATCH --time=00:01:00
#SBATCH --error=hello.out
#SBATCH --output=hello.out

module load goolf/1.6.10

mpirun -np 16 ./hello
```

Exercise 1 - Submit job to queue

- Submit the job to the queue

```
sbatch hello.sub
```

- Check the queue status

```
squeue
```

- Check the contents of the output file

```
cat hello.out
```

Exercise 2 - CUDA devices accessible by each process

- Load the goolf (g - GCC compiler, o - OpenMPI, o - OpenBLAS, l - ScaLAPACK, f - FFTW) and CUDA modules:

```
module load goolf
```

```
module load CUDA
```

- Compile the code

```
nvcc -lmpi mpi-cuda.cu -o mpi-cuda
```

Exercise 2 - Create job script

```
#!/bin/bash

#SBATCH --job-name=mpi-cuda
#SBATCH --nodes=1
#SBATCH --ntasks-per-node=8
#SBATCH --time=00:01:00
#SBATCH --gres=gpu:2
#SBATCH --error=mpi-cuda.out
#SBATCH --output=mpi-cuda.out

module load CUDA/5.5.22
module load goolf/1.6.10

mpirun -np 8 ./mpi-cuda
```

Exercise 2 - Submit job to queue

- Submit the job to the queue

```
sbatch mpi-cuda.sub
```

- Check the queue status

```
squeue
```

- Check the contents of the output file

```
cat mpi-cuda.out
```


Thank you!

Questions?