

Assignment 4: Using MPI on the CDS cluster

John Wallin

March 10, 2010

1 Introduction

This document was created to help you learn how to use MPI on the cds workstations. You should run a simple MPI code on this system. Once you have these steps completed, you can start developing and using MPI codes on these machines.

You will also turn in public rsa file to me so we can establish accounts on the larger gmice machine.

2 Setting up your SSH keys

1. Backup your old .ssh folder, if it exists.

```
CDS01:~>cp -R .ssh .ssh-old
```

2. If the directory doesn't exist, create a new .ssh directory. cd into this directory.

```
CDS01:~>mkdir .ssh  
CDS01:~>cd .ssh
```

3. Generate a new rsa key pair using the ssh-keygen command. Make sure to use a non-trivial passphrase.

```
CDS01:~>ssh-keygen  
Generating public/private rsa key pair.
```

```

Enter file in which to save the key (/Users/jwallin/.ssh/id_rsa):
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /Users/jwallin/.ssh/id_rsa.
Your public key has been saved in /Users/jwallin/.ssh/id_rsa.pub.
The key fingerprint is:
7b:36:8f:13:3b:10:d0:b5:6e:ad:bd:bd:2e:7d:f7:d9 jwallin@cds01
The key's randomart image is:
+--[ RSA 2048]-----+
|      . . .      |
|      . . .      |
|      . . .      |
|      . . .      |
|      S.o .      |
|      .o.o       |
|      ..=o..     |
|      o++o. =    |
|      .ooo++E|   |
+-----+

```

4. Concatenate the new private key into an identity file.

```
CDS01:~> cat id_rsa.pub > authorized_keys
```

5. Test the key out

```

CDS01:~> ssh cds02
The authenticity of host 'cds02 (129.174.112.89)' can't be established.
RSA key fingerprint is 93:ed:12:2a:59:3c:3b:80:eb:30:67:df:e0:16:17:aa.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added 'cds02,129.174.112.89' (RSA) to the list of known hosts.
Enter passphrase for key '/Users/jwallin/.ssh/id_rsa':
Last login: Thu Mar  4 17:38:00 2010 from cds01.gmu.edu
...
CDS02:~> exit

```

6. Once this stage is working, copy your public file into a new file called “email.pub” where “email” is your official Mason

email account. Send this file to me via email with the subject line "csi 702 hw 4".

```
cp id_rsa.pub email.pub
```

then download and email to me....

7. Put the passphrase into RAM

```
CDS01:~> eval 'ssh-agent'
Agent pid 23340
CDS01:~> ssh-add
Enter passphrase for /Users/jwallin/.ssh/id_rsa:
Identity added: /Users/jwallin/.ssh/id_rsa (/Users/jwallin/.ssh/id_rsa)
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
@                WARNING: UNPROTECTED PRIVATE KEY FILE!                @
@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
Permissions 0644 for '/Users/jwallin/.ssh/identity' are too open.
It is recommended that your private key files are NOT accessible by others
This private key will be ignored.
jwallin@CDS02:~>
```

8. Test the new phasephrase out

```
CDS01:~> ssh cds02
Last login: Wed Mar 10 02:32:51 2010 from cds01.gmu.edu

=====
Welcome to the CDS computer lab
=====

..
> exit
```

9. Repeat the logins for all the computers you need to reach for your mpi job

```
CDS01:~> ssh cds03
...
CDS03:~> exit
CDS01:~> ssh cds04
...
CDS04:~> exit
.
.
.
```

3 Compiling and running an MPI code

1. Find and download an MPI code. My favorite starter codes are pi3.f and cpi.c, and both can be found with a quick web search. Copy them into the top level of your home directory. Compile them using mpif90 and mpicc

```
CDS01:~> mpif90 pi3.f -o pi3
CDS01:~> mpicc cpi.c -o cpi
```

2. Create a host file for running on MPI. This should be a simple textfile named “hostfile” (or something similar) that contains the names of the machines you want to run the MPI program on. To start, just try using the local workstation. Then add more. The file should look like:

```
CDS01:~> cat hostfile
cds01.gmu.edu
cds02.gmu.edu
cds03.gmu.edu
CDS01:~>
```

3. You can now run the code. Use the command:

```
CDS01:~> mpirun --hostfile hostfile -np 6 --mca btl tcp,self ./pi3
```

The “-np 6” means run six copies of the code. Obviously, this can be changed. The “-mca btl tcp,self” is needed to let the mpirun command know that these workstations don’t use infiniband networking cards.