



USER MANUAL

Supercomputing Research & Education Center (ScREC)
Research Center for Modeling & Simulation (RCMS)
National University of Sciences & Technology (NUST)

v 1.1
2016

Table of Contents

1. Introduction	2
2. Introduction to ScREC	2
2.1 System Overview	3
2.2 Storage Area Network (SAN).....	5
2.3 Quick Summary of the System	5
2.4 Difference b/w Head Node and Compute Node.....	6
3. Using <i>afrit.rcms.nust.edu.pk</i> Supercomputer – First Steps.....	6
3.1 Requesting an account	7
3.2 Accessing <i>afrit.rcms.nust.edu.pk</i>	7
3.2.1 Accessing from Windows	7
3.2.2 Accessing from Linux.....	11
3.2.3 Home Directory	12
4. Password Management	12
5. Transferring Files	13
5.1 Transferring Files from Windows	13
5.2 Transferring Files from Linux.....	14
5.2.1 Command Line.....	14
5.2.2 Linux gFTP Applicaton	15
6. Accessing Compute Nodes	16
7. Accessing Software on Supercomputer.....	16
7.1 Setting Environment Variables.....	17
8. Running Jobs	18
8.1 Submitting Serial Jobs	18
8.2 Submitting Parallel Jobs	19
9. Contact Us.....	21

1. Introduction

This manual is intended to provide the minimum amount of information needed by a new user of afrit.rcms.nust.edu.pk Supercomputer. As such, it assumes that the user is familiar with many of the standard features of Supercomputer and Linux. Here you can find most of the information you need to use this supercomputing facility. Please read this manual carefully and if any doubt arises, do not hesitate to [contact us](#).

2. Introduction to ScREC

Supercomputing Research and Education Center (ScREC) is a center of excellence in High Performance Computing with a mission to support research and development efforts at National University of Sciences and Technology (NUST). It was established in 2012 at Research Center for Modeling and Simulation (RCMS). It houses the fastest supercomputer (afrit.rcms.nust.edu.pk) in academia in Pakistan till date. The facility is being utilized in computation-intensive research projects in areas of Fluid Dynamics, Biosciences, Huge Data Processing Applications such as Flood and Weather Forecasting, Financial Analysis, Oil and Gas Exploration, Energy Efficient Building Designs and Transportation Management at national-level.



Figure 1: afrit.rcms.nust.edu.pk

2.1 System Overview

Afrit.rcms.nust.edu.pk is installed in state of art Data Center with 80 KVA of UPS backup and 12-ton precision cooling system. The Data Center is protected by FM-200 based Automatic Fire Detection and Suppression System and manual Fire Extinguishers. CCTV Cameras and Access Control systems are being procured for effective surveillance of the Data Centre.

Afrit.rcms.nust.edu.pk is based on HP ProLiant DL380 & HP ProLiant DL160se G6 servers with Linux Operating System, an Infiniband high-speed interconnection, a high performance and reliable Storage Area Network (SAN) and NVidia Tesla GPUs.

See below a summary of the system:

- **02 Head Nodes.** Each one is equipped with:
 - Two 2.27 GHz 64bit Intel 4-core Xeon E5520 processors
 - 8 physical cores (16 logical cores if using Hyper-Threading)
 - 16GB DDR3 RAM
 - 2 x 450GB SATA Hard Drives
 - 10GigE connection to the Internet
 - 40Gbps QRD InfiniBand high-speed interconnection for internal communication
 - CentOS 6.5 Operating System
- **32 Compute Nodes.** Each one is equipped with:
 - Two 2.27 GHz 64bit Intel 4-core Xeon E5520 processors
 - 8 physical cores (16 logical cores if using Hyper-Threading)
 - 24GB DDR3 RAM
 - 2 x 250GB SATA Hard Drives
 - 40Gbps QRD InfiniBand high-speed interconnection for internal communication
 - CentOS 6.5 Operating System
 - NVidia Tesla S1070 GPU
 - 4 x 1 Tesla T10 Processors
 - CUDA Driver Version/ Runtime Version: 6.5
 - CUDA Capability Version: 1.3
 - 4 x 240 CUDA Cores
 - GPU Clock Rate of 1.44GHz

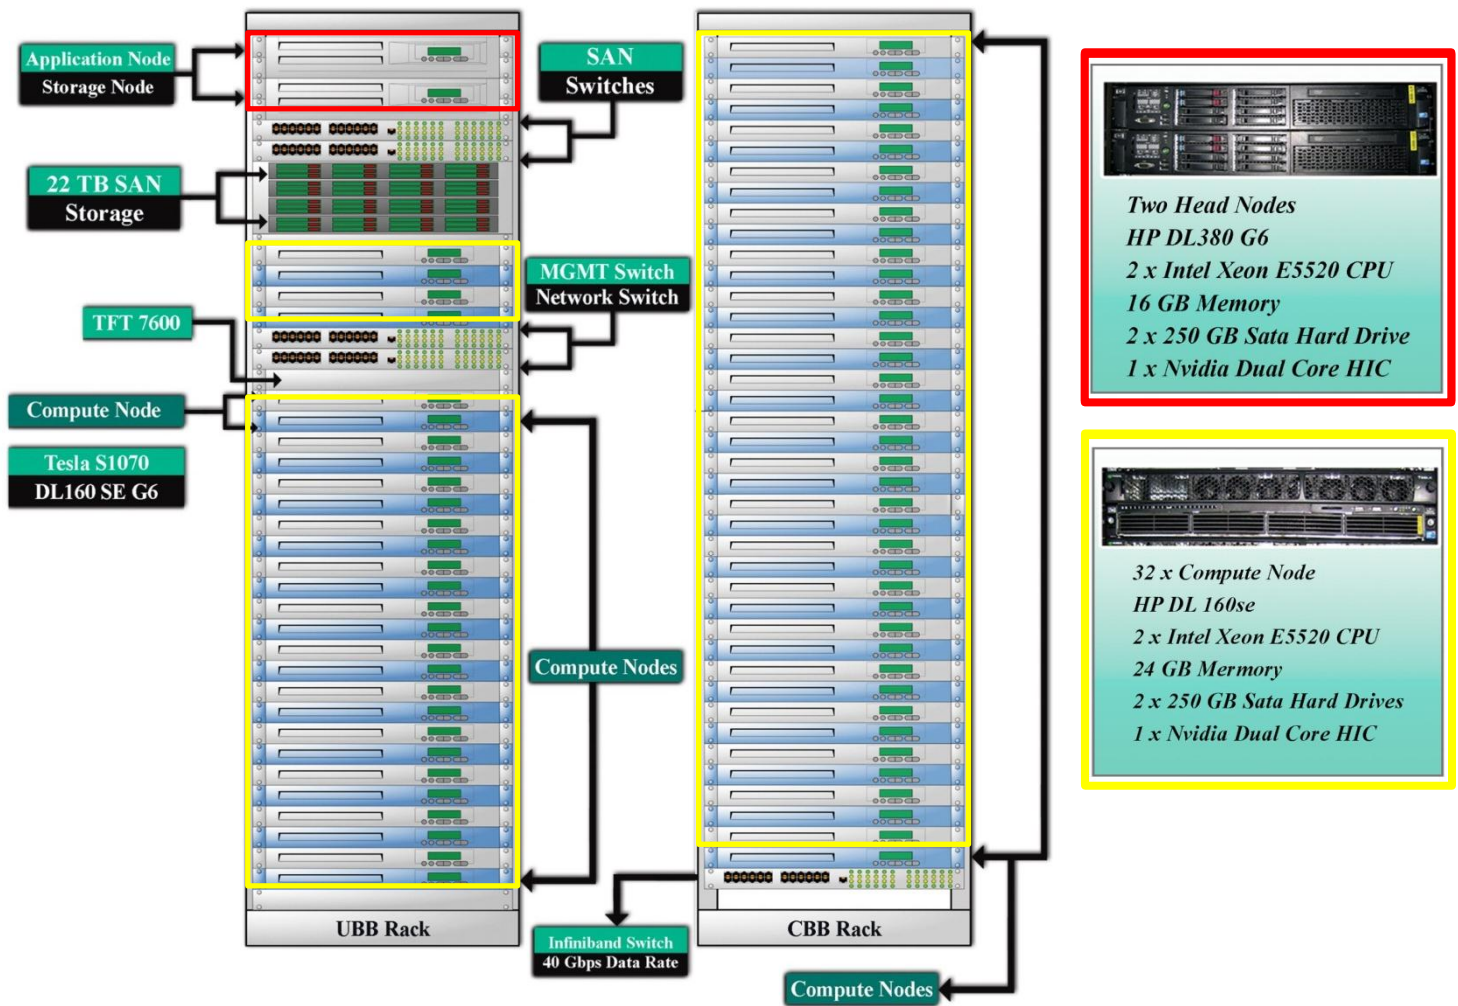


Figure 2: Logical Diagram of Supercomputer

2.2 Storage Area Network (SAN)

Total 22TB of SAN Storage is available for storing users' data. Two SAN Switches are installed in UBB Rack. 8 x 8 Gb transceiver is installed in each of the SAN Switch. Total 48 slots are available with storage capacity of 450GB/ slot. The system is configured on RAID-1 one unit and RAID-5 on four units, each containing 16 drives. One online spare drive is marked in each of the disk enclosure for high availability. In case of drive failure, the online spare drive will take over and data will be re-created depending on the RAID level.



Figure 3: SAN Setup

2.3 Quick Summary of the System

Host Name	afrit.rcms.nust.edu.pk
Total Number of Nodes	34
Number of Processing Cores	272
Total Memory	1.312TB
Operating System	CentOS Linux 6.5
Interconnects	Infiniband Switch
Storage	22TB
Total Number of GPU Cores	30720
Peak Performance	132 Teraflops

Table 1: Quick Summary of the System

2.4 Difference b/w Head Node and Compute Node

Typically, when you access a Supercomputing system, you are accessing its head node or login node. A head node is often nothing more than a simply configured system that is configured to act as a middle point between the actual cluster and the outside network. From a best practices standpoint, head nodes are often setup to provide a point of access and testing of the programs only. **Ideally one should not run computational intensive programs on the head node itself.**

On the other hand, the compute nodes are the workhorse of the cluster. One should use compute nodes for computational intensive work both serial and parallel, either in batch mode or interactive mode. A list of all available compute nodes is given below. Please read [section 6](#) of this manual to learn how to access compute nodes.

Compute Nodes				
compute-0-0	compute-0-6	compute-0-12	compute-0-18	compute-0-24
compute-0-1	compute-0-7	compute-0-13	compute-0-19	compute-0-25
compute-0-2	compute-0-8	compute-0-14	compute-0-20	compute-0-26
compute-0-3	compute-0-9	compute-0-15	compute-0-21	compute-0-27
compute-0-4	compute-0-10	compute-0-16	compute-0-22	compute-0-28
compute-0-5	compute-0-11	compute-0-17	compute-0-23	

Table 2: Available Compute Nodes

3. Using afrit.rcms.nust.edu.pk Supercomputer – First Steps

In order to make use of this facility, you will first need an account. In order to obtain an account, please follow the instructions in the “Requesting an Account” section below. Although this guide assumes that you have an account and have received your login credentials, prospective users who do not as yet have an account will hopefully find this manual useful for developing an understanding of how to use this facility. Once your account is created, we suggest you to read sections 3 to 8 of this user manual which describes how to login to the system, how to set user environment and how to submit jobs to the queue for execution.

3.1 Requesting an account

Follow the link below and download the “Supercomputer Account Application Form”. Fill-up this form and forward to RCMS. You will be informed about your supercomputer login credentials as soon as your account is created. You can access supercomputer once you get your username and password.

- <http://www.nust.edu.pk/INSTITUTIONS/Centers/RCMS/AboutUs/facilities/screc/Pages/Downloads.aspx>

3.2 Accessing *afrit.rcms.nust.edu.pk*

First thing you should know is your username and password. Once you have a login ID and its associated password you can get into the cluster via secure shell, SSH a remote login program, from anywhere in the world. The detailed method of access will depend on whether you are seeking to connect from a Linux system or a Windows system. Each case is described below.

3.2.1 Accessing from Windows

On a Windows system, you need to download and install following software before accessing supercomputer:

- PuTTY - <http://www.filehippo.com/search?q=putty>
- WinSCP - <http://www.filehippo.com/search?q=winscp>
- Xming - <http://sourceforge.net/projects/xming/files/Xming/6.9.0.31/>
- Xming Fonts - <http://sourceforge.net/projects/xming/files/Xming-fonts/7.5.0.93/>

After installing all of the above mentioned software, run and configure PuTTY. Depending on your location, enter one of the following hostname & port number in marked fields 1, 2 & 3 as shown in figure 4.

S#	Hostname/ IP Address	Port No	Location
1	10.9.41.100	22	Inside NUST H-12 Campus
2	External IP*	2299	Outside NUST H-12 Campus

Table 3: Hostname & Port No

It is necessary to configure PuTTY to use X-server program (Xming) on your windows machine if you are intended to use an application that has/ uses a graphical user interface. To do that, expand SSH in category section and click X11. You'll see a check box under X11

*External IP will be provided on request

forwarding. Mark this box to enable X11 forwarding. Next, write **localhost:0** in X display location. See below figure 5.

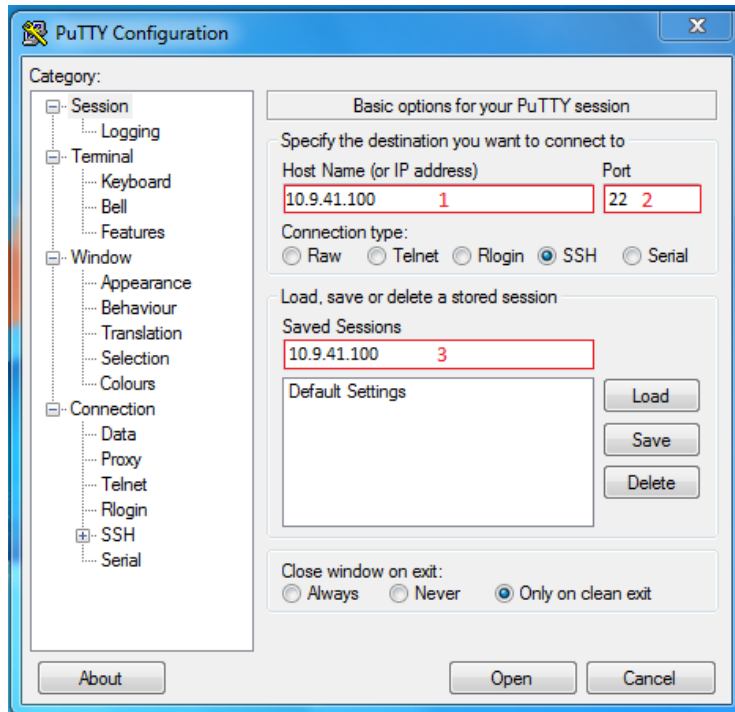


Figure 4: Entering Hostname in PuTTY

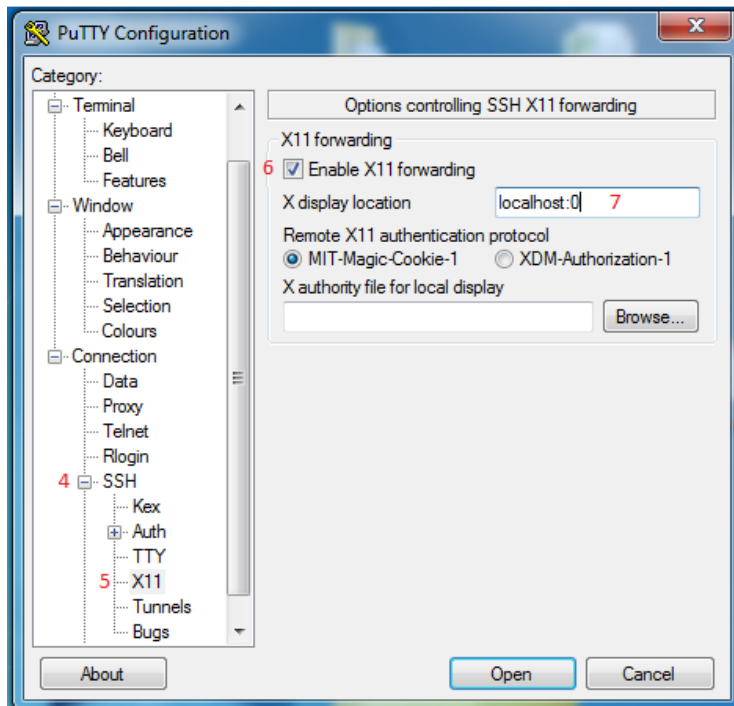


Figure 5: X11 Settings in PuTTY

After that, click on “Session” in category section and save these settings by clicking save button. These settings will be saved under “Default Settings” box. See below figure 6.

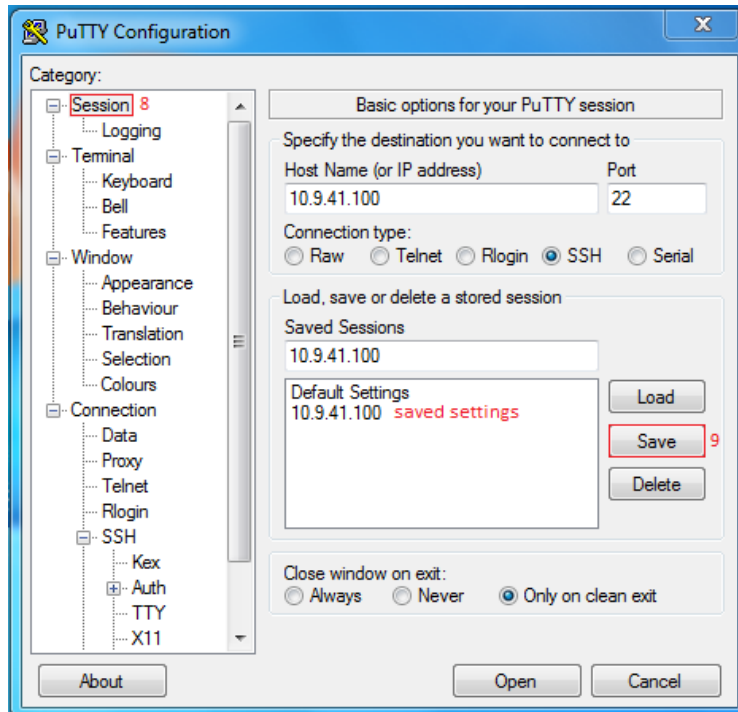


Figure 6: Saving Settings in PuTTY

After that select the settings you just saved and click open as shown in figure 7.

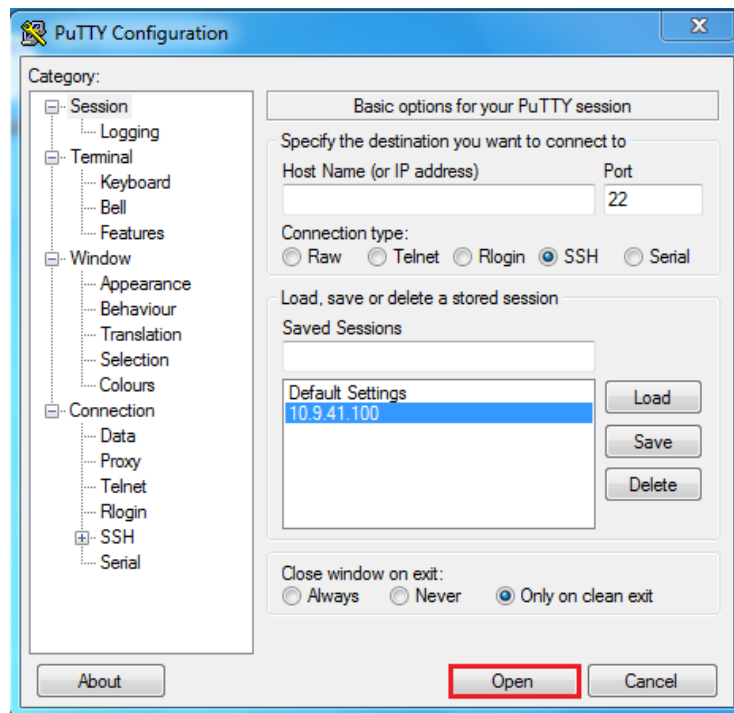


Figure 7: Connecting to Supercomputer

If it is your first time login to the supercomputer, you will get a warning message indicating that the host key from the server is unknown and will ask you if you agree to cache the new host key. Click Yes.

You'll see a black window as shown below. Type your username and password to login:


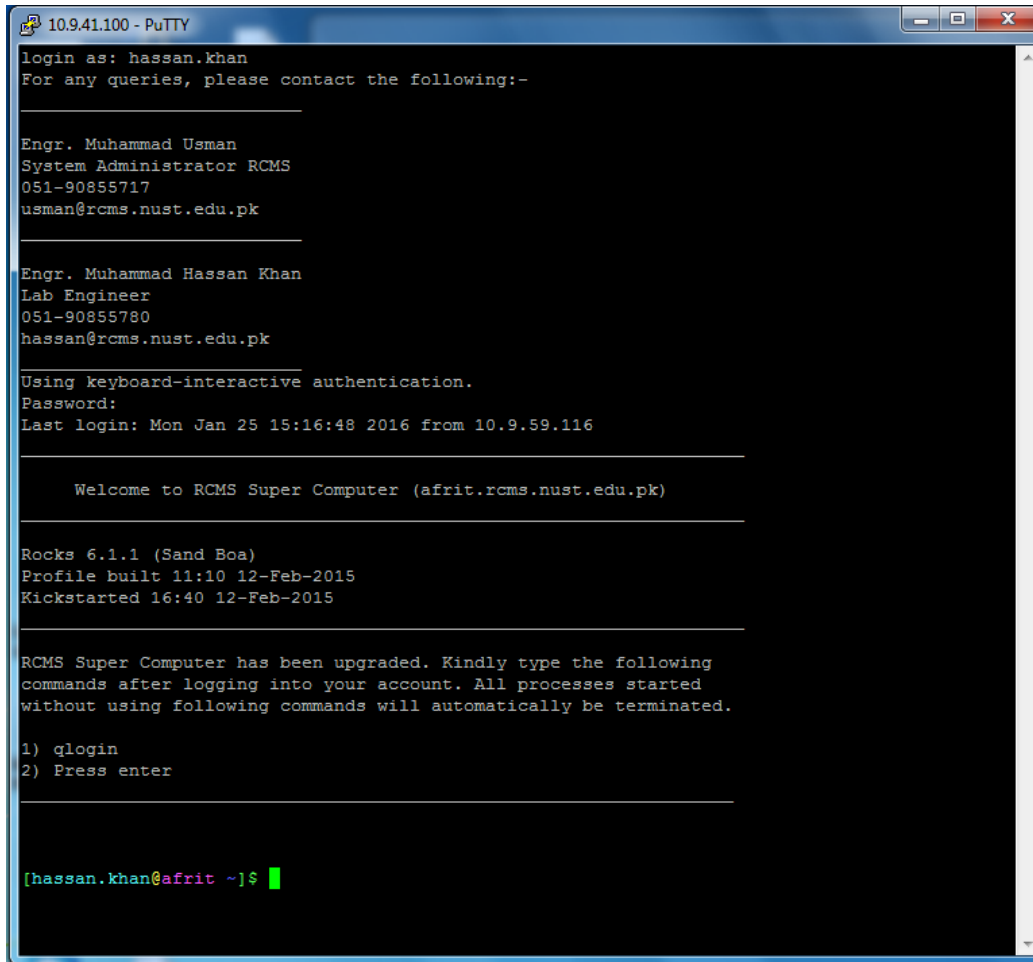


Figure 8: PuTTY Login Screen

Make sure that you have valid username and password to access supercomputer. In case you get any error message, please check your internet connection and Hostname in PuTTY session. Please [contact us](#) if error pertains.

On successful login, your PuTTY session will look something like as shown in figure 9.



```
10.9.41.100 - PuTTY
login as: hassan.khan
For any queries, please contact the following:-

-----
Engr. Muhammad Usman
System Administrator RCMS
051-90855717
usman@rcms.nust.edu.pk

-----
Engr. Muhammad Hassan Khan
Lab Engineer
051-90855780
hassan@rcms.nust.edu.pk

-----
Using keyboard-interactive authentication.
Password:
Last login: Mon Jan 25 15:16:48 2016 from 10.9.59.116

-----
Welcome to RCMS Super Computer (afrit.rcms.nust.edu.pk)

-----
Rocks 6.1.1 (Sand Boa)
Profile built 11:10 12-Feb-2015
Kickstarted 16:40 12-Feb-2015

-----
RCMS Super Computer has been upgraded. Kindly type the following
commands after logging into your account. All processes started
without using following commands will automatically be terminated.

1) qllogin
2) Press enter

-----
[hassan.khan@afrit ~]$
```

Figure 9: Successful Login to Supercomputer

3.2.2 Accessing from Linux

afrit.rcms.nust.edu.pk can be accessed using the `ssh` command on a Linux machine. Open Linux terminal and use one of the following commands to access supercomputer:

- `ssh -X username@10.9.41.100` (inside NUST H-12 campus)
- `ssh -p 2299 -X username@ExternalIP*` (outside NUST H-12 campus)

where “username” should be replaced by your username. Note that the `-X` option enables X11 forwarding – this will be required if you are intending to run graphical applications on the Supercomputer.

For example, if your username is hassan.khan, then you should use:

- `ssh -X hassan.khan@10.9.41.100` (inside NUST H-12 campus)
- `ssh -p 2299 -X hassan.khan@ExternalIP*` (outside NUST H-12 campus)

*External IP will be provided on request

3.2.3 Home Directory

A home directory, also called a login directory, is the directory on UNIX like operating systems that a user is first in after logging into the system. It is created automatically when a user account is created. The name of a user's home directory is by default identical to that of the user ID. Thus, for example, a user with a user name *hassan.khan* would typically have a home directory named *hassan.khan*. It is used to store user's data. A security restriction has been implemented on user's home directory to maintain the privacy of user's data. Users can only access their home directory. Type following command to find the exact path of your home directory:

```
echo $HOME
```

It will return something like */export/home3/rcms/hassan.khan*

A default quota will be enforced on all users to limit the amount of data stored here. It is highly discouraged to store unnecessary data on supercomputer.

4. Password Management

It is recommended that you should change your password on your first login. Your password should contain a mix of lower and upper case letters and numbers.

Use *passwd* command to change password.

```
[hassan.khan@afrit ~]$ passwd
Changing password for user hassan.khan.
Changing password for hassan.khan.
(current) UNIX password:
New password:
Retype new password:
passwd: all authentication tokens updated successfully.
```

Please note that your password will not be displayed while typing for security reasons.

5. Transferring Files

The detailed method of file transfer will depend on whether you are connecting from a Linux system or a Windows system. Each case is described below.

5.1 Transferring Files from Windows

On a Windows system, use WinSCP to transfer files from/ to the supercomputer. Run WinSCP and type appropriate hostname, port no, username and password and click login as shown in figure 10.




Figure 10: WinSCP Settings

Click continue if any authentication banner or message appears. A new window will appear containing two sides. The left side shows your “My Documents” folder on your local machine. The right side shows your home directory on the supercomputer. You can transfer files from/ to cluster account by simply using drag and drop function. Drag files from one side and drop them on the other side to upload or download them. You can also transfer files by selecting them and clicking the copy button. See figure 11.

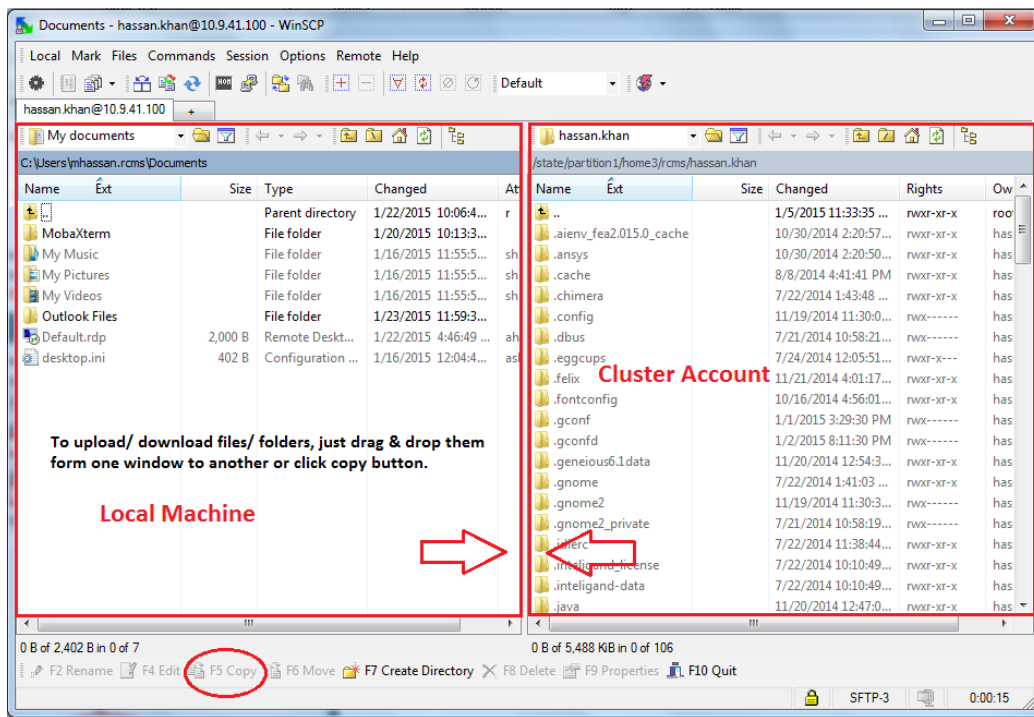


Figure 11: Transfer Files using WinSCP

5.2 Transferring Files from Linux

There are two ways you can use to transfer your files from/ to Supercomputer under Linux environment.

5.2.1 Command Line

You can use Secure Copy program to transfer your data to/ from supercomputer and local machine. Follow the below mentioned command syntax to transfer your data to supercomputer account:

scp /path/to/source/file username@hostname:/path/to/destination/

Follow the below mentioned command syntax to download your data from supercomputer account:

scp username@hostname:/path/to/source/ /path/to/destination/

Replace username with your login ID and hostname with one of the IP address as mentioned in [section 3.2.1](#).

Please use -rp switch with scp command if you get any errors. One such possible error is **filename is not a regular file**.

scp -rp /path/to/source/file username@hostname:/path/to/destination/

`scp -rp username@hostname:/path/to/source/ /path/to/destination/`

Type `man scp` on your terminal or follow below mentioned link for more details about scp command.

<https://www.garron.me/en/articles/scp.html>

5.2.2 Linux gFTP Applicaton

The other way to transfer your data to/ from supercomputer is to use gFTP program. Type following command in your terminal to install gFTP program.

`sudo apt-get install gftp` (Ubuntu only)

Once this software is installed, type `gftp` command to open it.

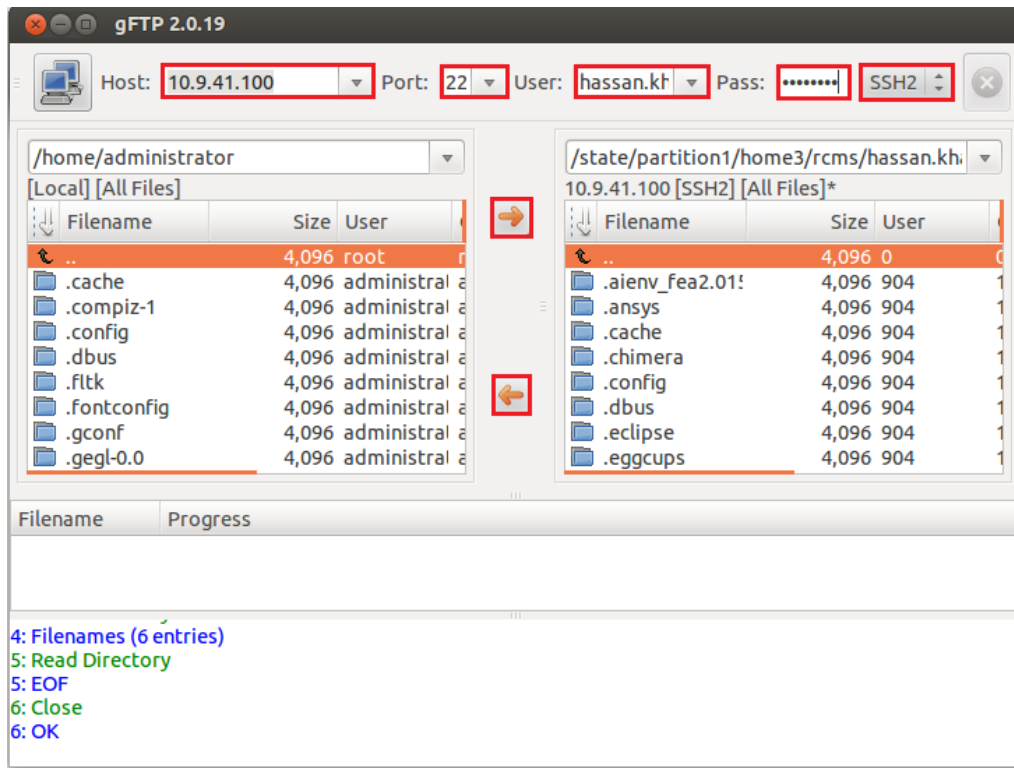


Figure 12: gFTP Settings

To connect to cluster, fill in the highlighted fields as shown in figure 12 and press enter. Once you are connected, transferring files is as simple as selecting the ones you want to upload or download and clicking the arrows as shown in figure 12. You can also transfer files by dragging and dropping them into the opposite window. Please make sure that hostname and port no is appropriate.

6. Accessing Compute Nodes

It is assumed that you have read [section 2.4](#) of this manual and completely understands the differences between head node and compute node. Accessing compute nodes is very simple. Once logged in, you can use following command to access any compute node:

```
ssh name-of-compute-node
```

Replace *name-of-compute-node* with the actual name of compute node. See [table 2](#). For instance, if you want to access compute-0-10, you should use above command as:

```
ssh compute-0-10
```

Your home directory will be visible on all compute nodes that means you will be able to access your data on all compute nodes.

7. Accessing Software on Supercomputer

[Afrit.rcms.nust.edu.pk](http://afrit.rcms.nust.edu.pk) Supercomputer comprises a number of software to support the research activity at NUST. All these software are installed in a specific directory which is accessible by all supercomputer users. Users can access these software by setting environment variable(s) to their startup script files (**.bashrc & .bash_profile**). These files are usually hidden and are placed in home directory of every user. Users can view hidden files by using following command:

```
ls -a
```

These files are created by default when a user account is created. Users can also create these files using text editor (vi/ vim) if these files are not present. Follow the link below to download the list of all available software with their environment variable(s):

<http://www.nust.edu.pk/INSTITUTIONS/Centers/RCMS/AboutUs/facilities/screc/Pages/Downloads.aspx>

Users should add the environment variable(s) of software of interest to their startup script files. Once added, use following commands or re-login to source the variables:

```
source ~/.bashrc
```

```
source ~/.bash_profile
```

7.1 Setting Environment Variables

It is assumed that you have downloaded the list of all available software and interested in using CHARMM. As mentioned in the list, copy the environment path of CHARMM.

```
export PATH=$PATH:/share/apps/packages/charmm/c40b1/exec/gnu_M
```

Login to Supercomputer using your login credentials and type following commands to set the path to your startup script files:

1. vim **.bashrc**
2. Press **shift+g** to go to the last line
3. Press **o** to enable insert mode
4. Insert the copied path by clicking the right mouse button
5. Press **Esc** to enable command mode
6. Type **:x** and press **Enter** to save and quit the file
7. vim **.bash_profile**
8. Repeat steps 2-6

After inserting the relevant path(s), re-login to your account to source the variables.

Now, type appropriate commands to run CHARMM. If you don't know the commands, navigate to the directory you have just added. In this case, the directory is:

```
/share/apps/packages/charmm/c40b1/exec/gnu_M
```

Type **cd /share/apps/packages/charmm/c40b1/exec/gnu_M** to change to this directory.

Then type **ls** to list all commands. To learn more about these commands, download the manual/ user guide from the relevant software website.

8. Running Jobs

In this section, we discuss the process for running jobs on afrit.rcms.nust.edu.pk supercomputer. Make sure that you have added environment variable(s) of software of interest to your startup script files as mentioned earlier. You can either login to any compute node and start running your software binaries or you can submit your job to Sun Grid Engine (SGE). It is a high performance computing cluster software which is responsible for accepting, scheduling, dispatching and managing user jobs. You can submit jobs to SGE via shell scripts.

8.1 Submitting Serial Jobs

Here is an example of a serial job script. It basically executes simple Linux commands. You can replace these commands with your software commands.

```
#!/bin/bash
#
#$ -cwd
#$ -j y
#$ -S /bin/bash
#
date
echo AH! My First Serial job is running 😊
whoami
```

Entries starting with `#$` are treated as SGE options.

- `-cwd` means to execute the job for the current working directory
- `-j y` means to merge the standard error stream into the standard output stream. It means a single file will contain both errors & output. Do not use this if you want separate error and output files.
- `-S /bin/bash` specifies the interpreting shell for this job

To submit this serial job script, you should use `qsub` command as follows:

```
qsub -V your_script.sh
```

You'll get following message when you submit your job.

```
your job N ("your_script.sh") has been submitted
```

where N is the job ID of your job.

Enter the following command to retrieve status information about your job(s).

```
qstat
```

If you need to delete an already submitted job, you can use following command with your job ID

```
qdel N
```

where N is the job ID of your job.

8.2 Submitting Parallel Jobs

Running parallel jobs on afrit.rcms.nust.edu.pk requires users to have basic understanding of Message Passing Interface (MPI) and its implementations. Follow the link below to learn about MPI:

<https://computing.llnl.gov/tutorials/mpi/>

Before running parallel jobs, you have to do following tasks:

1. Check if your application supports parallel execution
2. Which MPI implementation is compatible with your application

Currently, following MPI implementations are installed on supercomputer:

1. Open MPI-1.8.4

```
export PATH=/usr/mpi/gcc/openmpi-1.8.4/bin:$PATH
export LD_LIBRARY_PATH=/usr/mpi/gcc/openmpi-1.8.4/lib64:$LD_LIBRARY_PATH
```
2. Open MPI-1.6.2 **(Default)**

```
export PATH=/opt/openmpi/bin:$PATH
export LD_LIBRARY_PATH=/opt/openmpi/lib:$LD_LIBRARY_PATH
```
3. MPICH2-1.4.1p1

```
export PATH=/opt/mpich2/gnu/bin:$PATH
export LD_LIBRARY_PATH=/opt/mpich2/gnu/lib:$LD_LIBRARY_PATH
```
4. MVAPICH2-2.0

```
export PATH=/usr/lib64/mvapich/bin:$PATH
export LD_LIBRARY_PATH=/usr/lib64/mvapich/lib:$LD_LIBRARY_PATH
```

Open MPI-1.6.2 is the default MPI but you can also use other MPI implementation. To use other MPI implementation, add environment variable of that implementation to your startup script files and source them.

Here is an example of parallel job script. You can use this script to run your simulation on as many nodes as you want.

```
#!/bin/bash
#
#$ -cwd
#$ -j y
#$ -S /bin/bash
#
mpirun -np N -hostfile machines your_commands input_file &
```

where:

- **N** is the numeric digit that specifies the number of processors you want to use i.e. 100
- **machines** is the text file containing the names of the compute nodes on which your job should run. You can create this file using vim/vi editor. You can calculate how many nodes you need to use 100 cores. Follow the link below to download sample machine file

The command to submit a MPI parallel job script is similar to submitting a serial job script.

```
qsub -V your_script.sh
```

You'll get following message when you submit your job.

```
your job N ("your_script.sh") has been submitted
```

where N is the job ID of your job.

Enter the following command to retrieve status information about your job(s).

```
qstat
```

If you need to delete an already submitted job, you can use following command with your job ID

```
qdel N
```

where N is the job ID of your job.

9. Contact Us

Please feel free to contact us if you have any query or need assistance.

S#	Name	Designation	Email ID	Contact No
1	Engr. Muhammad Usman	Sys Admin	usman@rcms.nust.edu.pk	051-9085-5717
2	Engr. Muhammad Hassan Khan	Lab Engr.	hassan@rcms.nust.edu.pk	051-9085-5780
3	Mr. Shahzad Shoukat	Asst. Sys Admin	shahzad@rcms.nust.edu.pk	051-9085-5714

Our address is:

Supper Computing Lab, 1st Floor, Acad1 Block
Research Center for Modeling and Simulation (RCMS)
National University of Sciences & Technology (NUST)
Sector H 12, Islamabad
Pakistan