HPC cluster Setup using ROCKS

Introduction

Computer clusters

The aim of the following exercises is to familiarize the user with the steps involved in creating and installing a 2 node HPC cluster based on the ROCKS - Open Source Toolkit for Real and Virtual Clusters (http://www.rocksclusters.org/)

The physical infrastructure of the cluster is 2 computers with a network cable between them.

- The ROCKS DVD may be downloaded from http://www.rocksclusters.org/
- After installation, you can use a web-browser to view the documentation on the local server using the URL http://localhost/
- Documentations and references may be found at

  http://www.rocksclusters.org/roll-documentation/

Exercise 1

Installing a Front-end/cluster Master node

Requirements

- ROCKS dvd
- Computer with 2 network cards

Steps

- Boot the machine and start the installation from the Rocks DVD, it may be necessary to press F12 or F9 or a similar key or even to configure BIOS to boot from DVD drive before hard-disk
- At the ROCKs screen, press a key quickly at the boot screen to get the boot: prompt
- Type **"build"**
- Click the button titled "CD/DVD-based Roll"
- Select all listed rolls and click submit
- Click Next after verifying that the selected Rolls are listed on the left had side of the screen
- Fill-in the various entries on the screen and use a hostname such as master.hpc and use a two-letter code for the country.
- Configure the public interface with DHCP or with a static IP address 192.168.10.10/255.255.255.0 and click Next
- For the private internal network, it should propose the 10.1.1.0/16 network with master as 10.1.1.1, Click Next
- You may need to fill-in the gateway and DNS entries, you can enter the IP used for themaster.hpc if necessary.
- Set a suitable root password
- Select the right time zone
- Use automatic repartitioning (note all existing data on the disks will be lost)
- The installation begins..
- After the server reboots, login as root
- Click System --> Administration --> Users and Groups to create a new user
- Logout as root and login as the user.
- Open a web browser by clicking the firefox icon on the menubar and type in the URL below to see information about your cluster
    - http://localhost
    - URL for local documentation is
        - http://localhost/roll-documentation
    - URL for monitoring
        - http://localhost/ganglia/

Exercise 2

Installing the cluster worker/compute nodes

Requirements

- Working ROCKS frontend/masternode
- Computer/server with single Ethernet

Steps

- Attached the new computer to the switch and also attach the eth0 interface of the masternode to the switch as well. If you do not have a switch, please connect the node directly to the masternode using an Ethernet cable.
- Open a terminal and become root user using the su command
    - su
- On the master node run the command insert-ethers in order to capture the MAC addresses of the compute node.
    - Insert-ethers
- On the menu select compute
- Now power on the first compute node, make sure you set it up to boot from the network, you may have to modify BIOS to enable
    - Network Card Option LAN
    - Network Boot
- The node should be captured and wait until it reports kickstart. The installation should start accordingly. You can monitor this using the rocks-console command tool

Exercise 3

Using SGE for job submission

Requirements

- Working ROCKS cluster (masternode and compute-0-0)

Steps

- Login as a normal user
- Create a job script with a suitable text editor (pico or nano) containing the following lines

```
#!/bin/bash
#
#$ -cwd
#$ -j y
#$ -S /bin/bash
#
date
/usr/bin/openssl speed
date
```

Option	Explanations
-cwd	Run in the current working directory
-j y	Stdout and stderr in the same output file
-S /bin/bash	Use the bash shell for running the job
-M *{email-address}*	Send notifications about job to email-address
-o {*filename}*	Send output into file *{filename}*
-e *{filename}*	Send stderr into file *{filename}*

- Save the file as test.sh
- Submit the job using the qsub command

  qsub test.sh

  Note: you can submit the test.sh job multiple times, repeat the qsub command about 5 times.

- Check the status of your job using the command qstat

qstat

qstat -f

- Delete one of the jobs using the qdel command

  qdel

- Other commands are qconf for checking which queues are available

  qconf –sql

  qmod –d to disable a queue and qmod –e to enable a queue

The SGE roll documentation roll-sge-usersguide.pdf contains an example for a parallel job.