



Building a Rocks Cluster

Rocks-A-Palooza II Lab Session

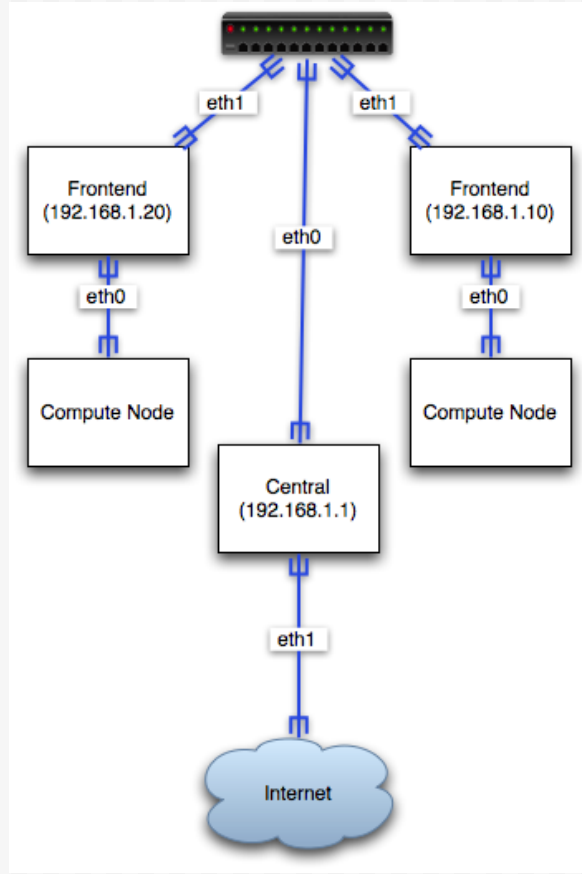
Cluster Building Time

- ◆ Break into Groups
- ◆ Every Group Grab
 - ⇒ 2 Servers
 - ⇒ 2 Power Cords
 - ⇒ 2 Ethernet Cables
 - 1 long
 - 1 short
 - ⇒ 1 Keyboard / Mouse
 - ⇒ 1 Monitor
- ◆ Small Clusters
 - ⇒ 1 frontend
 - ⇒ 1 compute
 - ⇒ 1 cross-over Ethernet cable (no switch)

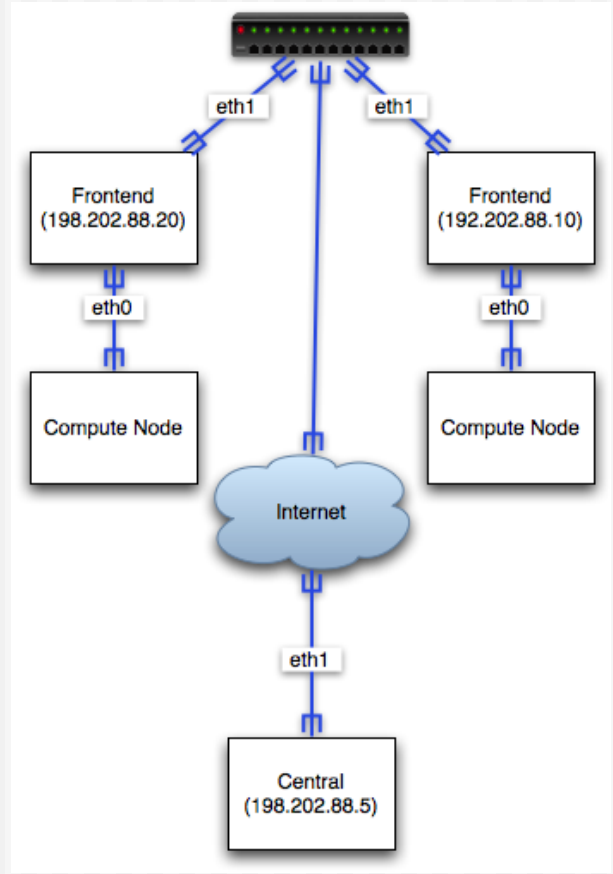




Today's Lab Network



lab



reality

Network Information

◆ Frontend Addresses

- ⇒ 192.168.1.10
- ⇒ 192.168.1.20
- ⇒ 192.168.1.30
- ⇒ 192.168.1.40
- ⇒ 192.168.1.50
- ⇒ 192.168.1.60
- ⇒ 192.168.1.70
- ⇒ 192.168.1.80

IP Address	192.168.1.xx
Netmask	255.255.255.0
Gateway	192.168.1.1
Nameserver	198.202.75.26

Start Installing Your Frontend

- ◆ Installation Methods
 - ⇒ CDs
 - ⇒ Central
- ◆ CD
 - ⇒ Slow
 - ⇒ Does not require a network
 - ⇒ Type frontend
 - ⇒ Then add all CD rolls
- ◆ Central
 - ⇒ Fast
 - ⇒ Requires a network
 - ⇒ Type frontend central=192.168.1.1

Rocks Cluster Distribution

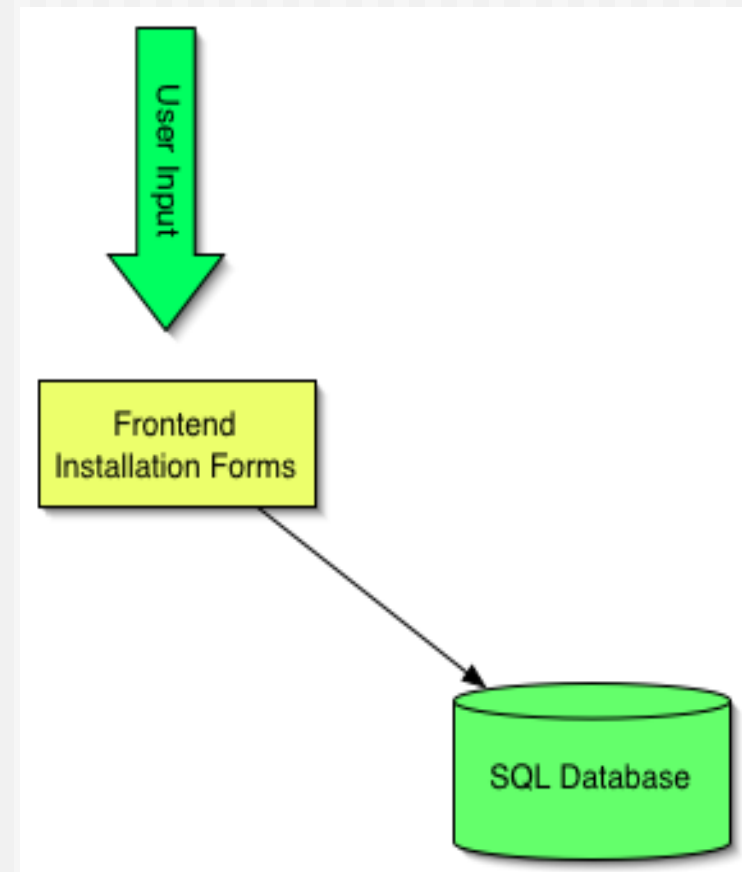
What do you want to kickstart?

- Frontend:
type "frontend"
- Upgrade your frontend:
type "frontend upgrade"
- Frontend Network Install
type "frontend central=name"
where name is "Rocks", or the
FQDN of your central server.
- Rescue
type "frontend rescue"
- Cluster node:
do nothing or press return


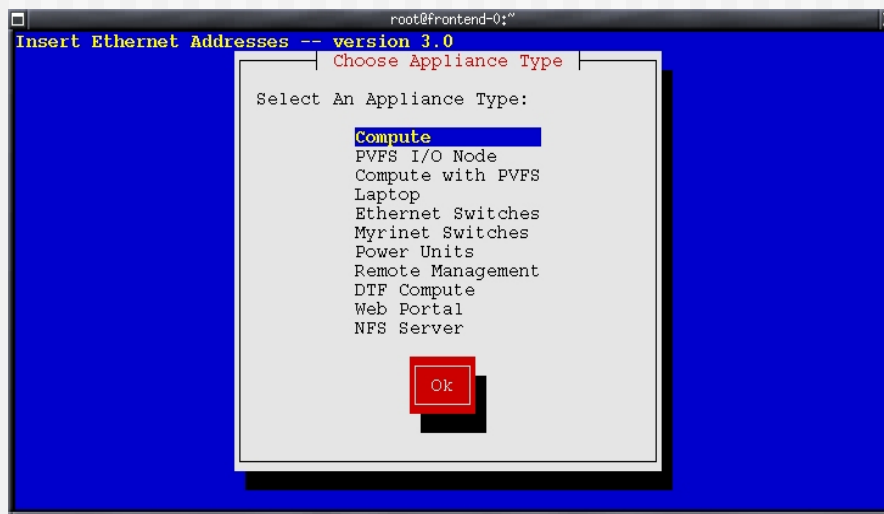


Interactive Screen

- ◆ Fill out the install 'screens'
- ◆ Use the provided network information
- ◆ Choose your own password
- ◆ All information goes into the cluster database



Add Compute Node with Insert-ethers



- ◆ Collect the Ethernet MAC address of cluster nodes
- ◆ Only done once, during integration
- ◆ Populates cluster database



Open Lab

- ◆ Rocks-A-Palooza
 - ⇒ Is about you guys
 - ⇒ Other topics
 - ⇒ Questions
- ◆ Adult Swim
 - ⇒ Go nuts on your clusters
 - ⇒ Globus
 - ⇒ SGE
 - ⇒ Configuration Graph



Frontend Installation

- ◆ Turn on node
- ◆ Insert 'Compute' CD
- ◆ At 'boot:' prompt,
type:

frontend central=192.168.1.1

Rocks Cluster Distribution

What do you want to kickstart?

- Frontend:
type "frontend"
- Upgrade your frontend:
type "frontend upgrade"
- Frontend Network Install
type "frontend central=name"
where name is "Rocks", or the
FQDN of your central server.
- Rescue
type "frontend rescue"
- Cluster node:
do nothing or press return



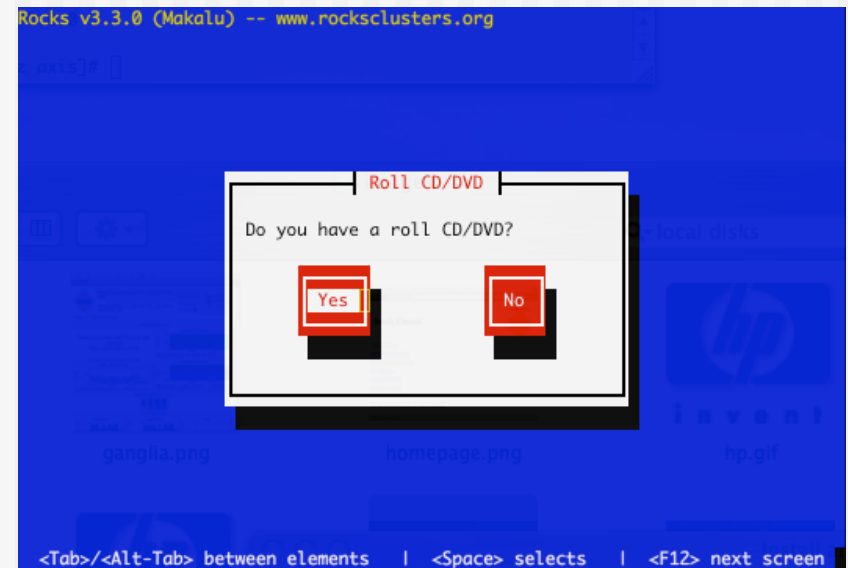


Rolls

- ◆ Anaconda Starts
- ◆ Discovers Rolls on 'central' server
- ◆ Select all rolls

Rolls

- ◆ Asked if have anymore Roll servers
 - ⇒ Select 'No'
- ◆ Asked if have any Roll CD/DVD media
 - ⇒ This is where you can add a roll that is not on a central server
 - ⇒ For this lab, select 'No'



Cluster Information

- ◆ Specific to Rocks
- ◆ Used for Certificates
 - ⇒ SSL/HTTPS
 - ⇒ Globus
- ◆ Hostname
 - ⇒ Must be FQDN
 - ⇒ Must be in DNS
 - ⇒ Must not be an Alias

Rocks v3.3.0 (Makalu) -- www.rocksclusters.org

Cluster Information

Fill in at least the FQDN

Fully Qualified Hostname: Country:

Cluster Name: Contact:

Organization: URL:

Locality: LatLong:

State:

OK Back

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

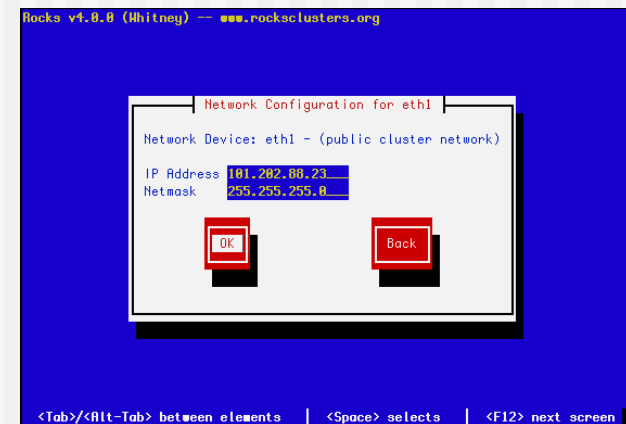

Partitioning

- ◆ Automatic
 - ➔ 6GB /
 - ➔ 1GB swap
 - ➔ Remainder for /export
- ◆ Manual
 - ➔ You choose
 - ➔ Must create a /export
- ◆ Select Wisely




Networks

- ◆ Private Network
 - ⇒ eth0
 - ⇒ Cluster-side only
- ◆ Public Network
 - ⇒ eth1
 - ⇒ Internet/LAN side
- ◆ You must configure both and have 2 NICs



Gateway

- ◆ Gateway / DNS
 - ➔ Same as any other device on your network
- ◆ All traffic for compute nodes is NATed through the frontend.
- ◆ DNS is only for the frontend, compute nodes use the frontend as their DNS.



Network Time Protocol

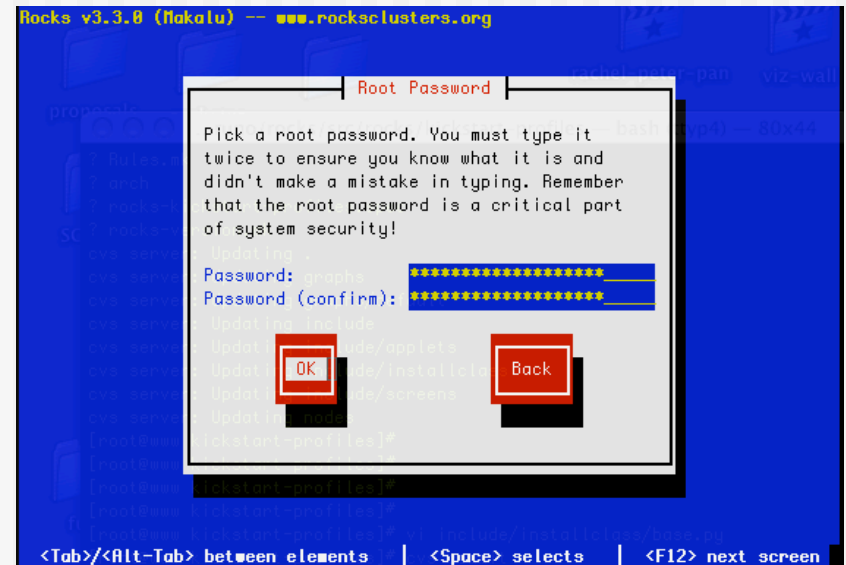
- ◆ Choose timezone
 - ⇒ UTC is a good choice
 - ⇒ Or localize
- ◆ Default NTP server is
 - ⇒ pool.ntp.org
 - ⇒ Change it if you wish



```
Rocks v3.3.0 (Makalu) -- www.rockclusters.org
Time Configuration
What time zone are you located in?
[*] System clock uses UTC
America/Montevideo
America/Montreal
America/Montserrat
America/Nassau
America/New_York
Which network time server would you like to use?
time.apple.com
OK Back
```


Root Password

- ◆ Password is secure
 - ⇒ Not stored in clear text form anywhere (not in DB)
- ◆ Also used for mysql password
- ◆ Also used for wordpress password
 - ⇒ When you want to add content to frontend's homepage
 - Which we'll do in the 'Cluster Management and Maintenance Lab'





Installing Packages

Rocks v3.3.0 (Makalu) -- www.rocksclusters.org

In (13)

Package Installation

Name : perl-5.8.0-88.7-i386
Size : 38104k
Summary: The Perl programming language.

99%


	Packages	Bytes	Time
Total :	708	2747M	0:09:15
Completed:	431	1450M	0:04:53
Remaining:	277	1297M	0:04:22

52%

<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

Integrate Compute Nodes

- ◆ Log into Frontend (as root)
- ◆ Run `insert-ethers`
 - ⇒ Can choose appliance type
 - ⇒ Rolls add new appliance types
 - ⇒ For now we will use Compute
- ◆ Turn on first node
 - ⇒ Nodes are integrated serially
 - ⇒ Need to map machine name to machine location
 - ⇒ After we integrate machines can be re-installed in parallel
- ◆ Remote Terminal (ekv)
 - ⇒ `ssh compute-0-0 -p2200`





Discovering Compute-0-0

```
Insert Ethernet Addresses -- version 3.2.0
Opened kickstart access to 10.0.0.0/255.0.0.0 network

Inserted Appliances

Press <F10> to quit, press <F11> to force quit
```

```
Insert Ethernet Addresses -- version 3.2.0
Opened kickstart access to 10.0.0.0/255.0.0.0 network

Inserted Appliances
00:30:c1:a0:ac:25    compute-0-0    < > #

Press <F10> to quit, press <F11> to force quit
```

```
Insert Ethernet Addresses -- version 3.2.0
Opened kickstart access to 10.0.0.0/255.0.0.0 network

Inserted Appliances
Discovered New Appliance
Discovered a new appliance with MAC (00:30:c1:a0:ac:25)

Press <F10> to quit, press <F11> to force quit
```

```
Inserted Appliances
00:30:c1:a0:ac:25    compute-0-0    (*) #

Retrieved kickstart file
```

useradd

```
root@rocks-39:~ — bash (tty1)
[root@rocks-39 ~]# useradd mjk
Creating user: mjk
make: Entering directory `/var/411'
/opt/rocks/sbin/411put --comment="#" /etc/auto.home
411 Wrote: /etc/411.d/etc.auto.home
Size: 579/253 bytes (encrypted/plain)
Alert: sent on channel 255.255.255.255:8649 with master 10.1.1.1

/opt/rocks/sbin/411put --comment="#" /etc/passwd
411 Wrote: /etc/411.d/etc.passwd
Size: 2816/1905 bytes (encrypted/plain)
Alert: sent on channel 255.255.255.255:8649 with master 10.1.1.1

/opt/rocks/sbin/411put --comment="#" /etc/shadow
411 Wrote: /etc/411.d/etc.shadow
Size: 1961/1272 bytes (encrypted/plain)
Alert: sent on channel 255.255.255.255:8649 with master 10.1.1.1

/opt/rocks/sbin/411put --comment="#" /etc/group
411 Wrote: /etc/411.d/etc.group
Size: 1236/740 bytes (encrypted/plain)
Alert: sent on channel 255.255.255.255:8649 with master 10.1.1.1

make: Leaving directory `/var/411'
[root@rocks-39 ~]# passwd mjk
Changing password for user mjk.
New UNIX password:
BAD PASSWORD: it is based on a (reversed) dictionary word
Retype new UNIX password:
passwd: all authentication tokens updated successfully.
[root@rocks-39 ~]#
```

411 Distributes User Info

- ◆ In previous slide, added a password for ‘mjk’
 - ⇒ This password **is** immediately available on the frontend
- ◆ This password **is not** immediately available on the compute nodes
 - ⇒ User id files (/etc/passwd, /etc/shadow, etc.) are distributed to the compute nodes by 411 service
 - ⇒ 411 broadcast updates every hour
 - ⇒ Or, can force the update:
 - # make -C /var/411 force

user login

```
mjk@rocks-39:~ — bash (ttyp1)
$~>
$~> ssh concave.rocksclusters.org
mjk@concave.rocksclusters.org's password:
Last login: Mon May 16 19:50:09 2005 from client64-84.sdsc.edu
Rocks Frontend Node - Rocks-39 Cluster
Rocks 4.0.0 (Whitney)
Profile built 13:03 26-Apr-2005

Kickstarted 13:03 26-Apr-2005

It doesn't appear that you have set up your ssh key.
This process will make the files:
    /home/mjk/.ssh/id_rsa.pub
    /home/mjk/.ssh/id_rsa
    /home/mjk/.ssh/authorized_keys

Generating public/private rsa key pair.
Enter file in which to save the key (/home/mjk/.ssh/id_rsa):
Created directory '/home/mjk/.ssh'.
Enter passphrase (empty for no passphrase):
Enter same passphrase again:
Your identification has been saved in /home/mjk/.ssh/id_rsa.
Your public key has been saved in /home/mjk/.ssh/id_rsa.pub.
The key fingerprint is:
17:44:24:f3:b7:bd:41:48:4a:82:83:a6:d1:5f:68:af mjk@rocks-39.sdsc.edu
[mjk@rocks-39 ~]$
```