

Kvm Users Guide



6.1 Edition



Kvm Users Guide :

6.1 Edition

Published Mar 06 2013

Copyright © 2013 University of California

This document is subject to the Rocks® License (see Appendix A: Rocks Copyright).

Table of Contents

Preface	v
1. Overview	1
2. Installing	2
2.1. On a New Server	2
2.2. On an Existing Server.....	2
3. Using the KVM Roll	3
3.1. Installing a VM Server	3
3.2. Installing VM Containers	3
3.3. Installing Virtual Clusters.....	4
3.4. Physical Frontend with Virtual Compute Nodes.....	13
3.5. Remotely Controlling VMs with Pilot	15
3.6. KVM Networking (Advanced).....	18
4. Command Reference	20
4.1. No Commands Installed at BUILD	20
A. Rocks® Copyright	21
B. Third Party Copyrights and Licenses	22
C. Known issues	23

List of Tables

1-1. Summary.....	1
1-2. Compatibility.....	1
3-1. Virtual Machine Names.....	5

Preface

The KVM Roll installs and configures Virtual Machines (VMs) on Rocks Clusters.

A physical frontend can configure VMs on client nodes (*VM container* appliances). A VM container is a physical machine that houses and runs VMs.

The KVM Roll also supports building virtual clusters. The frontend can be installed as a *VM server* appliance and the client nodes can be installed as VM containers. Then a virtual frontend can be installed on the VM server while virtual compute nodes can be installed on the VM containers. All network traffic is encapsulated within a unique VLAN, that is, each virtual cluster has its own VLAN.

Please visit the KVM site¹ to learn more about their release and the individual software components.

Notes

1. <http://www.linux-kvm.org>

Chapter 1. Overview

Table 1-1. Summary

Name	service-pack
Version	6.1
Maintained By	Rocks Group
Architecture	x86_64
Compatible with Rocks®	6.1

The service-pack roll has the following requirements of other rolls. Compatibility with all known rolls is assured, and all known conflicts are listed. There is no assurance of compatibility with third-party rolls.

Table 1-2. Compatibility

Requires	Conflicts
Base	
Kernel	
OS	
	Xen




This roll has been released independent of the corresponding Rocks® release. It therefore requires the complete **OS** roll and will not function correctly if using only the **Jumbo** or incomplete set of **OS** CDRoms.

Chapter 2. Installing

2.1. On a New Server

The service-pack roll should be installed during the initial installation of your server (or cluster). This procedure is documented in section 1.2 of the Rocks® usersguide. You should select the service-pack roll from the list of available rolls when you see a screen that is similar to the one below.



2.2. On an Existing Server

The service-pack Roll may also be added onto an existing server (or frontend). For sake of discussion, assume that you have an iso image of the roll called `service-pack.iso`. The following procedure will install the Roll, and after the server reboots the Roll should be fully installed and configured.

```
$ su - root
# rocks add roll service-pack.iso
# rocks enable roll service-pack
# cd /export/rocks/install
# rocks create distro
# rocks run roll service-pack | bash
# init 6
```

Chapter 3. Using the KVM Roll

3.1. Installing a VM Server

A VM Server is machine that can house virtual frontend appliances. It is required to build a VM Server if you wish to build virtual clusters.

Building a VM Server is just like building a traditional frontend, except that you *must have installed properly the KVM Roll*. Follow the procedure [Install and Configure Your Frontend¹](#) and be sure to supply the KVM Roll.

After you build the VM Server, you'll need to install VM Containers (see the next section).

3.2. Installing VM Containers

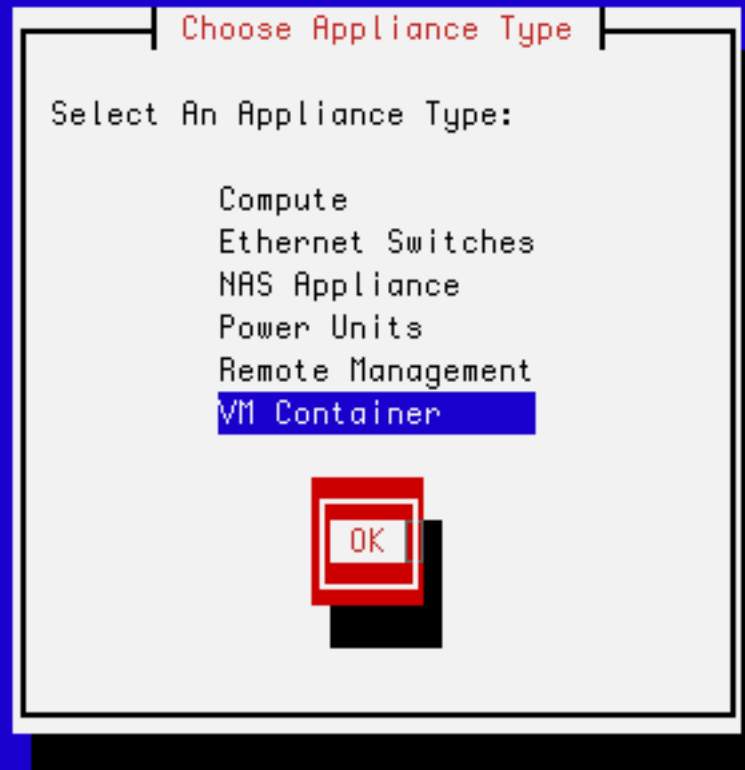
You will need to install a physical machine that will act as the "container" for your VMs. This method is very similar to the method for installing compute nodes.

On the frontend, execute:

```
# insert-ethers
```

You will see a screen that looks like:


```
Insert Ethernet Addresses -- version 5.0
Opened kickstart access to 10.0.0.0/255.0.0.0 network
```



Select the 'VM Container' appliance, then hit 'OK'.

Now PXE boot the physical machine that will be your VM container. Just like a compute node, the VM container will be recognized by insert-ethers and installed. The default name of the node will be `vm-container-X-Y`.

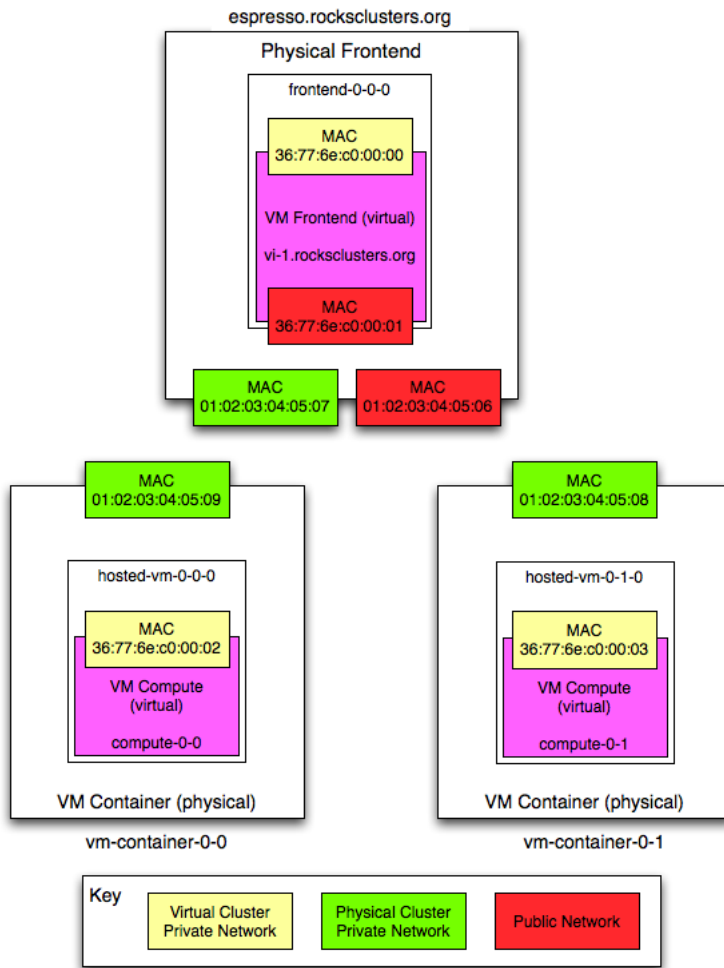
You can install as many VM containers as you like.

3.3. Installing Virtual Clusters

3.3.1. Provisioning a Virtual Cluster

After you install a VM Server and at least one VM Container, you are ready to provision a virtual cluster.

We'll use the following illustration as a guide to help keep track of the names of the physical machines and the virtual machines.



In the above picture, "espresso.rocksclusters.org" is a physical machine. Also, "vm-container-0-0" and "vm-container-0-1" are physical machines that were kickstarted by "espresso". The machine "frontend-0-0-0" is a virtual machine that is hosted by "espresso". The machines "hosted-vm-0-0-0" and "hosted-vm-0-1-0" are VMs that are associated with "frontend-0-0-0" (they are all in the same VLAN).

Depending on your perspective, the virtual machines have different names. Dom0 is a physical machine that hosts (multiple) virtual systems. DomU are guests and generally refer to names by usual convention. The equivalence is:

Table 3-1. Virtual Machine Names

Host	Dom0 Name (physical)	DomU Name (virtual)
37:77:6e:c0:00:00	frontend-0-0-0	vi-1.rocksclusters.org
37:77:6e:c0:00:01	hosted-vm-0-0-0	compute-0-0
37:77:6e:c0:00:02	hosted-vm-0-1-0	compute-0-1



An important point is that the only common thing between the physical side and the virtual side is the MAC address (in yellow). We will use the MAC address of a virtual machine to control it (e.g., to initially power it on).

The names in the virtual cluster look like the names in a traditional cluster -- the frontend is named "vi-1.rocksclusters.org" and its compute nodes are named "compute-0-0" and "compute-0-1". If you login to "vi-1.rocksclusters.org", you would be hard pressed to tell the difference between this virtual cluster and a traditional physical cluster.



You must select your own IP address for your virtual frontend. The IP address "137.110.119.118" is managed by UCSD and should not be used by you.

They are only used here to show you a concrete example.

First, we'll add a virtual cluster to the VM Server's database. In this example, we'll add a frontend with the IP of "137.110.119.118" and we'll associate 2 compute nodes with it:

```
# rocks add cluster ip="137.110.119.118" num-computes=2
```

The above command will take some time and then output something similar to:

```
created frontend VM named: frontend-0-0-0
created compute VM named: hosted-vm-0-0-0
created compute VM named: hosted-vm-0-1-0
```

The command adds entries to the database for the above nodes and establishes a VLAN that will be used for the private network (eth0 inside the VM).

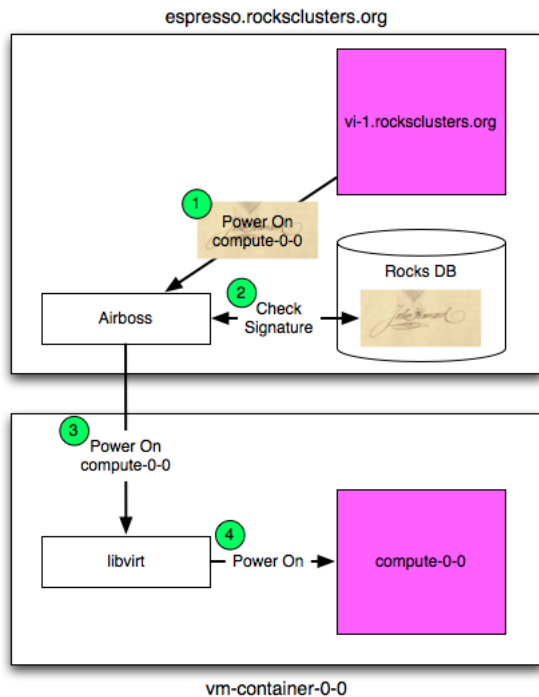
Info about all the defined clusters on the VM Server (including the physical cluster) can be obtained with the command: `rocks list cluster`:

```
# rocks list cluster
FRONTEND          CLIENT NODES      TYPE
espresso.rocksclusters.org: ----- physical
:                 vm-container-0-0 physical
:                 vm-container-0-1 physical
frontend-0-0-0-public: ----- VM
:                 hosted-vm-0-0-0  VM
:                 hosted-vm-0-1-0  VM
```

3.3.2. The Airboss

In Rocks, we've developed a service known as the "Airboss" that resides on the physical frontend (in Dom0) and it allows non-root users to control their VMs. The motivation for this service is that libvirt (a virtualization API written by RedHat that can control several different virtualization implementations) assumes "root" access to control and monitor VMs.

The Airboss in Rocks is a small service that uses digitally signed messages to give non-root users access to their virtual cluster (and only their virtual cluster). The Airboss relies upon public/private key pairs to validate messages. The administrator of the physical hosting cluster must issue a single command to associate a public key with a particular virtual cluster. At that point, the full process of booting and installing a virtual cluster can be controlled by the (authorized) non-root user.



In the above picture, a user that is logged in to `vi-1.rocksclusters.org` wants to power on `compute-0-0` (one of the VMs associated with the virtual cluster). The user executes the "power on" command. The command creates a "power on" message, signs it with a private key, then sends it to the Airboss that is running on `espresso.rocksclusters.org`. The Airboss verifies the message signature. If the signature is valid, then the Airboss instructs libvirt on `vm-container-0-0` to start ("power on") `compute-0-0`.

3.3.3. Creating an RSA Key Pair

Before we can install a VM, we must create an RSA key pair. These keys will be used to authenticate Airboss commands. To create a key pair, execute:

```
# rocks create keys key=private.key
```

The above command will ask for a pass phrase for the private key. If you would like a "passphraseless" private key, execute:

```
# rocks create keys key=private.key passphrase=no
```

The above command will place your private key into the file `private.key` and it will output the public key for your private key:

```
# rocks create keys key=private.key
Generating RSA private key, 1024 bit long modulus
.....+++++
.....+++++
e is 65537 (0x10001)
Enter pass phrase for private.key:
Verifying - Enter pass phrase for private.key:
Enter pass phrase for private.key:
writing RSA key
-----BEGIN PUBLIC KEY-----
MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDMoCPmR/Kev64znRBxvtsniXIF
```

```
dyQMxR/bBFKNDmvmzPuPUim5jmD3TLilnH75/KidtJCwlb+Lhr5Cs6/9sRzX6rX2
ExVUZsgo4A+O+XMk8KeowO/c2rPc+YdXaBir3Aesm/MCfCZaidZae8QLmVKW7Va5
qEr19gyhhR7uDX+hgwIDAQAB
-----END PUBLIC KEY-----
```

Now save the public key to file, that is, copy the above public key:

```
-----BEGIN PUBLIC KEY-----
MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDMoCPmR/Kev64znRBxvtsniXIF
dyQMxR/bBFKNDmvmzPuPUim5jmD3TLilnH75/KidtJCwlb+Lhr5Cs6/9sRzX6rX2
ExVUZsgo4A+O+XMk8KeowO/c2rPc+YdXaBir3Aesm/MCfCZaidZae8QLmVKW7Va5
qEr19gyhhR7uDX+hgwIDAQAB
-----END PUBLIC KEY-----
```

And save your public key into a file (e.g., \$HOME/public.key).

We now want to associate your public key with the virtual cluster you provisioned. This will allow you to use your private key to send authenticated commands to control your cluster. To associate your public key with your virtual cluster, execute:

```
# rocks add host key frontend-0-0-0 key=public.key
```

We can see the relationship by executing:

```
# rocks list host key
HOST          ID PUBLIC KEY
frontend-0-0-0: 7  -----BEGIN PUBLIC KEY-----
:                MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDMoCPmR/Kev64znRBxvtsniXIF
:                dyQMxR/bBFKNDmvmzPuPUim5jmD3TLilnH75/KidtJCwlb+Lhr5Cs6/9sRzX6rX2
:                ExVUZsgo4A+O+XMk8KeowO/c2rPc+YdXaBir3Aesm/MCfCZaidZae8QLmVKW7Va5
:                qEr19gyhhR7uDX+hgwIDAQAB
:                -----END PUBLIC KEY-----
:                -----
```

We see that the public key is associated with "frontend-0-0-0" (the name of the VM in Dom0).

3.3.4. Installing a VM Frontend

Now, we'll want to install the virtual frontend. First, login to the physical frontend (e.g., espresso). To start the VM frontend install, we'll need to power on and install the VM frontend:

```
# rocks set host power frontend-0-0-0 action=install key=private.key
```




The action of "install" ensures that the VM will be put into install mode, then it will be powered on.

Then, to connect to the VM's console, execute:

```
# rocks open host console frontend-0-0-0 key=private.key
```

Soon you will see the familiar frontend installation screen:



In the "Hostname of Roll Server" field, insert the FQDN of your VM Server (the name of the physical machine that is hosting the VM frontend). Then click "Download".

From here, you want to follow the standard procedure for bringing up a frontend² starting at Step 8.

After the VM frontend installs, it will reboot. After it reboots, login and then we'll begin installing VM compute nodes.

3.3.5. Installing VM Compute Nodes

Login to the VM frontend (the virtual machine named "vi-1.rocksclusters.org" in the example picture at the top of this page), and execute:

```
# insert-ethers
```

Select "Compute" as the appliance type.

In another terminal session on vi-1.rocksclusters.org, we'll need to set up the environment to send commands to the Airboss on the physical frontend. We'll do this by putting the RSA private key that we created in section Creating an RSA Key Pair (e.g., private.key) on vi-1.rocksclusters.org.

Prior to sending commands to the Airboss, we need to establish a ssh tunnel between the virtual frontend (e.g., vi-1) and the physical frontend (e.g., espresso, where the Airboss runs). This tunnel is used to securely pass Airboss messages. On the virtual frontend (e.g., vi-1), execute:

```
# ssh -f -N -L 8677:localhost:8677 espresso.rocksclusters.org
```

Now we can securely send messages to the Airboss.

Now, we're ready to install compute nodes. But, there's a problem - when we first login to vi-1.rocksclusters.org, the only machine we know about is ourself (i.e., vi-1.rocksclusters.org). There are no other nodes in the virtual

frontend's database. But the physical machine knows about the MAC addresses of the virtual compute nodes (e.g., hosted-vm-0-0-0 and hosted-vm-0-1-0) that are associated with this virtual cluster. The good news is, we can ask the Airboss on the physical frontend for a list of MAC addresses that are assigned to our virtual cluster:

```
# rocks list host macs vi-1.rocksclusters.org key=private.key
```

Which outputs:

```
MACS IN CLUSTER
36:77:6e:c0:00:02
36:77:6e:c0:00:00
36:77:6e:c0:00:03
```

The MAC address 36:77:6e:c0:00:00 is ourself (the VM frontend) and the other two MACs (36:77:6e:c0:00:02 and 36:77:6e:c0:00:03) are the VM compute nodes that are associated with our VM frontend.


We can use the MAC address of the VM compute nodes to power up and install our compute nodes:

```
# rocks set host power 36:77:6e:c0:00:02 key=private.key action=install
```



The action of "install" ensures that the VM will be put into install mode, then it will be powered on.

Soon, you should see insert-ethers discover the VM compute node:



After the virtual compute node is discovered by insert-ethers, we can open a console to the node by executing:

```
# rocks open host console compute-0-0 key=private.key
```

Lastly, to power off a virtual compute node (e.g., compute-0-0), execute:

```
# rocks set host power compute-0-0 key=private.key action=off
```


3.3.6. Using RedHat's Virt-Manager (Root Users Only)

Virt-manager is a program produced by RedHat that is a desktop user interface for managing virtual machines. This section describes how to use some of virt-manager's features to control and monitor VMs on a Rocks cluster.


To interact with the VM frontend's console, on the physical frontend, you need to start "virt-manager":

```
# virt-manager
```

This will display a screen similar to:



Double click on the "localhost" entry and then you'll see:

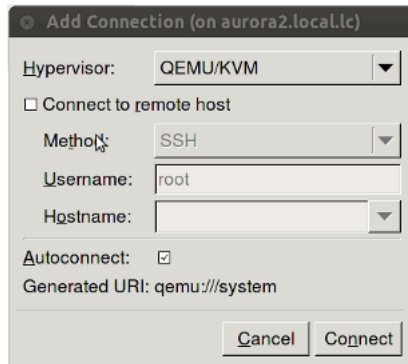


To bring the up the console for the VM frontend, double click on "frontend-0-0-0".


Now we'll describe how to connect to the console for the virtual compute node "compute-0-0". In the example

configuration described at the top of this page, the VM "compute-0-0" is hosted on the physical machine named "vm-container-0-0" so we'll need to tell "virt-manager" to open a connection to "vm-container-0-0".

Inside "virt-manager", click on "File" then "Open connection...". This brings up a window that looks like:




Now change the "Connection:" field to "Remote tunnel over SSH" and enter "vm-container-0-0" for the "Hostname:" field:




Then click "Connect".

In the "virt-manager" window, you should see something similar to:

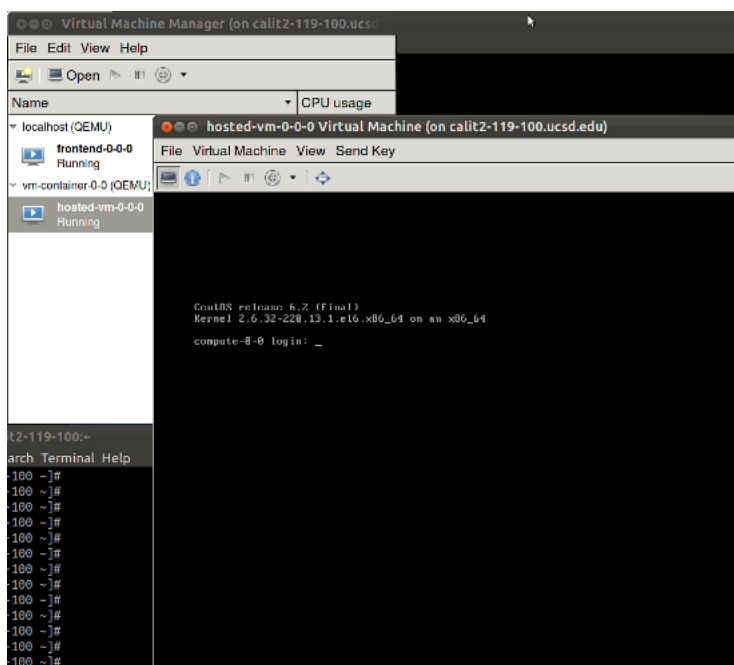


Double click on "vm-container-0-0" and then you'll see:



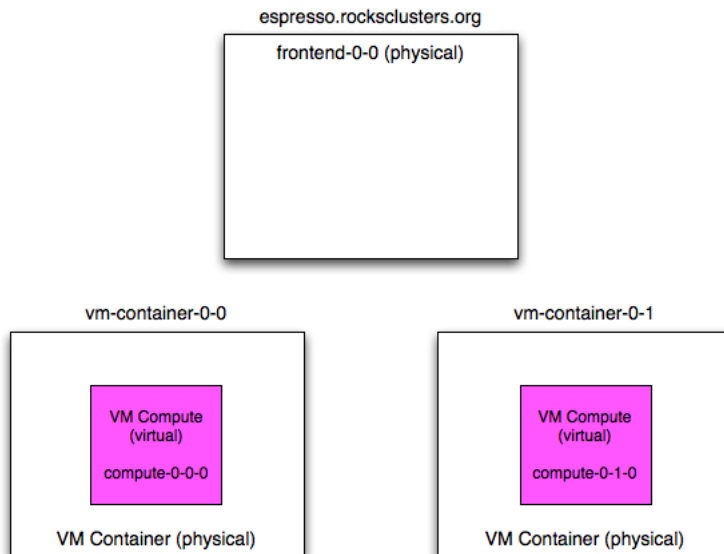
Now to connect to the compute node's console, double click on "hosted-vm-0-0-0". Recall that from the perspective of the physical frontend (the VM Server), "hosted-vm-0-0-0" is the name for the VM "compute-0-0" (again, see the figure at the top of this page).

You should now see the console for compute-0-0:



3.4. Physical Frontend with Virtual Compute Nodes

In this scenario, the frontend is a physical machine (not a VM) and the compute nodes are virtual machines.



In the above picture, "frontend-0-0" is a physical machine (with the public name of "espresso.rocksclusters.org"). The physical machine "frontend-0-0" controls two VM compute nodes named "compute-0-0-0" and "compute-0-1-0". This means that "compute-0-0-0" and "compute-0-1-0" are configured by "frontend-0-0". This is opposed to the "virtual cluster scenario" (Installing Virtual Clusters), where the virtual frontend ("frontend-0-0-0") configured the VM compute nodes, and the physical machine that housed "frontend-0-0-0" only started and stopped the virtual compute nodes.

3.4.1. Adding, Installing and Booting VMs with a Physical Frontend

In the common case, you will execute three Rocks commands over the lifetime of your VMs: add (to add VM info to the database), start (to boot or install a VM) and stop (to shutdown an installed VM).

To add a VM to the system, you need to associate a VM with a physical machine (i.e., a VM container) and you need to assign an appliance type to the VM. Here's an example:

```
# rocks add host vm vm-container-0-0 membership="Compute"
```

The above command will output a message similar to:

```
added VM compute-0-0-0 on physical node vm-container-0-0
```

This tells us that, in the database, the compute VM named "compute-0-0-0" has been assigned to the physical machine "vm-container-0-0".

The next step is to install the VM.

VMs are installed with the `start`. Here's how to install the VM that was added above:

```
# rocks start host vm compute-0-0-0
```



After running the command above, you may see the following message:

```
libvir: KVM Daemon error : POST operation failed: (kvm.err "Error creating domain: Disk isn't accessible")
```

This is not a problem. The above means the file that will be used for the VMs disk space was not present when the VM was started. The "rocks start host vm" command eventually creates it and starts the VM. That is, while you may see the error message above, the VM is actually running and installing.

The above command will start the standard Rocks installation process for the VM named "compute-0-0-0". After the installation process initializes the network inside the VM, you can monitor the installation just like a physical machine installation by executing:

```
# rocks-console compute-0-0-0
```

After the installation completes, the VM will reboot. After the VM boots, you can interact with the VM just like any other physical machine.

3.5. Remotely Controlling VMs with Pilot

3.5.1. Pilot Overview

"Pilot" is a program that can be used on Windows, Mac OS X and Linux systems to remotely and securely control VMs that are housed on Rocks Clusters. Pilot communicates to the "Airboss" (see section The Airboss for more info) on a Rocks frontend. Pilot can power off, power on, install and connect to the console of VMs.

3.5.2. Installing Pilot

3.5.2.1. Installing Pilot on Mac OS X

Open a "terminal" session and execute:

```
$ curl -o pilot.py http://www.rocksclusters.org/ftp-site/pub/rocks/extra/pilot/pilot.py
$ chmod a+x pilot.py
$ curl -o TightVncViewer.jar http://www.rocksclusters.org/ftp-site/pub/rocks/extra/pilot/osx/Tigh
$ curl -o foundation-python-extras.dmg http://www.rocksclusters.org/ftp-site/pub/rocks/extra/pilo
```

Install the "foundation-python-extras.dmg" package:

```
$ open foundation-python-extras.dmg
```

This will open a window with a picture of an open box icon with the label "foundation-python-extras.pkg" under it. Double click on the open box icon.

This will open an installation window. Click on the default settings to properly install this package.

3.5.2.2. Installing Pilot on Windows

The following procedure has been tested on Windows 7, 32-bit.

First, you'll need to install "curl" on your machine. Download and unzip the following file:

<http://am.net/lib/TOOLS/curl/curl-7.21.2-ssl-sspi-zlib-static-bin-w32.zip>

Then, from the Rocks web site, download:

<http://www.rocksclusters.org/ftp-site/pub/rocks/extra/pilot/windows/DLSupport.bat>

Open a Windows Command Prompt and execute:

```
PATH=%PATH%;<directory for unzipped curl>
DLSupport
```

The above will download: Pilot, VCredist, Python v2.7, OpenSSL and M2Crypto.

Install VCredist, Python v2.7, OpenSSL and M2Crypto by executing:

```
vcredist_x86.exe
python-2.7.msi
Win32OpenSSL_Light-1_0_0a.exe
M2Crypto-0.20.2.win32-py2.7.exe
```

If you don't already have an SSH client, you'll need to download one. We suggest Putty:

<http://the.earth.li/~sgtatham/putty/latest/x86/putty-0.60-installer.exe>

Most Windows PCs already have Java installed. Pilot has been tested with versions available from:

<http://www.java.com/en/download/>

3.5.3. Create a Public/Private Key Pair

Pilot uses a private/public key pair to authenticate its messages with the Airboss. If you haven't already created a key pair and associated the public key with the virtual frontend of the cluster you wish to remotely control, then login to the physical frontend that houses your virtual cluster and execute the procedure found here: [Creating an RSA Key Pair](#).

Then copy the private key to the same directory on your machine where you have installed "pilot.py".

3.5.4. Open an SSH Tunnel from Your Machine to the Physical Frontend

Pilot sends messages to the Airboss on port 8677. We'll need to open a secure tunnel and forward port 8677 from your local machine to the Rocks frontend where the Airboss is running.

3.5.4.1. Mac OS X SSH Tunneling

Open a new terminal session and execute:

```
$ ssh -L 8677:localhost:8677 root@beopen.rocksclusters.org
```

Be sure to replace "beopen.rocksclusters.org" with the FQDN of your frontend.

3.5.4.2. Windows SSH Tunneling

If using PUTTY, a guide to setting up tunnels can be found here:

http://docs.cs.byu.edu/general/ssh_tunnels.html#use-putty-to-set-up-a-tunnel

3.5.5. Using Pilot for Remote VM Power Control

Before you get started, open a new terminal session on your machine. Pilot is a command-line utility.

To power on a VM, you must know the MAC address of the VM you wish to power on. Assuming you want to power on the VM named "frontend-0-0-0" on your physical frontend, execute:

```
# rocks list host interface frontend-0-0-0
SUBNET  IFACE  MAC                IP                NETMASK          MODULE NAME          VLAN O
private eth0   76:77:6e:40:00:00  10.1.255.251     255.255.0.0       ----- frontend-0-0-0      2    -
public  eth1   76:77:6e:40:00:01  137.110.119.118  255.255.255.0     ----- frontend-0-0-0-public 0    -
```

The VM frontend-0-0-0 has the MAC addresses of 76:77:6e:40:00:00 and 76:77:6e:40:00:01. You can use either with pilot.

To power on frontend-0-0-0, execute:

```
$ ./pilot.py set host power host=76:77:6e:40:00:00 key=private.key action=on
```

To power off frontend-0-0-0, execute:

```
$ ./pilot.py set host power host=76:77:6e:40:00:00 key=private.key action=off
```

To power on and to force frontend-0-0-0 to install, execute:

```
$ ./pilot.py set host power host=76:77:6e:40:00:00 key=private.key action=install
```

3.5.6. Using Pilot to Connect to a VM's Console

You can also connect to a VM's console with pilot by executing:

```
$ ./pilot.py open host console host=76:77:6e:40:00:00 key=private.key
```



When you connect to a VM's console, you will see two mouse pointers, that is, the mouse pointers are not in "sync". You will have to experiment to find which pointer is the true pointer for the console.

3.6. KVM Networking (Advanced)

In order to support KVM Virtual Machines it is necessary to have in place a particular network configuration on the hosting servers (which can be frontends or vm-containers). In particular bridges must be set up to provide Virtual Machine with network connectivity (KVM networking differ substantially from XEN networking). In this paragraph we will explain in much details what are the differences from a standard Rocks Cluster network configuration.

3.6.1. VM Network Bridging to Physical Devices

When a VM is bridged to the physical device, it must be assigned in the same subnet as the physical device with a compatible IP address (this case is exactly what is described in Physical Frontend with Virtual Compute Nodes).

For this reason every Ethernet interfaces is renamed with a 'p' (for physical) in front of its original name. The interface is also put into promiscuous mode and then attached to a bridge named with the original name of the interface. The IP address and netmask are finally configured on the bridge device.

For example if you have an eth0 which is supposed to be configured with IP address 1.2.3.4 and netmask 255.255.255.0 you will end up with the physical interface named peth0 in promiscuous mode attached to a bridge named eth0 with IP address 1.2.3.4 and netmask 255.255.255.0. The creation of the bridge device, the renaming of the physical interface and the assignment of the IP address is all managed by Rocks through the Red Hat init scripts. This means that `rocks sync host network` creates the proper configuration files in `/etc/sysconfig/network-scripts/`.

If you start a virtual compute on a VM Container, with the VLAN 0 (the virtual machine belongs to the same cluster as the physical frontend), you will see a virtual device belonging to the virtual machine attached to the eth0 bridge.

```
[root@vm-container-0-0 ~]# brctl show
bridge name bridge id STP enabled interfaces
eth0 8000.00144f80de00 no peth0
      vnet0
virbr0 8000.5254007953c7 yes virbr0-nic
```

As you can see from the previous example bridge eth0 is connected to peth0 (the physical interface) and vnet0 (the virtual interface of the virtual machine). virbr0 is a default bridge started by libvirt which you can ignore.

3.6.2. Logical VLAN Devices

In this scenario, The guest (hosted-vm-0-0-0) and the host (vm-container-0-0) are not in the same logical network (this is the scenario described in Provisioning a Virtual Cluster).

The virtual machine hosted-vm-0-0-0 will be connected to the network through a Logical VLAN (for more information consult the Base Roll documentation on networking³). Logical Vlan interfaces don't need bridge device and they are dynamically loaded only when the virtual machine is started. They are called after the physical name interface simply adding the number of the VLAN (e.g. VLAN 2 on peth0 will be called peth0.2). Virtual machines are attached to the VLAN interface using a MacVTap⁴ driver which does not require any bridging infrastructure. After you run `rocks start host vm hosted-vm-0-0-0` if you go on the physical node you will see an interface called peth0.2 (the physical peth0 with the VLAN tag set to 2) and the macvtap0 which is the interface used by the virtual host.

Notes

1. </roll-documentation/base/6.1/install-frontend.html>
2. </roll-documentation/base/6.1/install-frontend.html>
3. <http://www.rocksclusters.org/roll-documentation/base/6.0/x1051.html#NETWORKING-LOGICAL-VLAN>
4. <http://virt.kernelnewbies.org/MacVTap>

Chapter 4. Command Reference

4.1. No Commands Installed at BUILD

install rocks-command-kvm

Appendix A. Rocks® Copyright

Appendix B. Third Party Copyrights and Licenses

This section enumerates the licenses from all the third party software components of this Roll. A "best effort" attempt has been made to insure the complete and current licenses are listed. In the case of errors or omissions please contact the maintainer of this Roll. For more information on the licenses of any components please consult with the original author(s) or see the Rocks® CVS repository¹.

Notes

1. <http://cvs.rocksclusters.org>

Appendix C. Known issues

This is the list of known issue with the KVM roll:

1. Ethernet channel bonding does not work with the KVM roll (rocks add host bonded).
2. Nvidia MCP77 integrated Ethernet Controller does not handle properly vlan tagged packages, so virtual clusters will not work on this hardware.
3. We have found that when a virtual frontend and a virtual compute node run on the same physical node it is necessary to disable tx checksumming offload on the eth0 in order to properly install the compute node. If during PXE boot compute node fails with "Connection timed out" error message run the following command on the frontend:

```
ethtool -K eth0 tx off
```

To make this change persistent after a reboot add that line in /etc/rc.d/rc.local of your frontend.

4. If you have a Frontend node and a compute node connected directly with a cable (and no switch in between), you have to change the STP parameters of the private network bridge otherwise the installation of virtual cluster nodes will fail. To do this run:

```
brctl setfd eth0 1.00  
brctl sethello eth0 0.5
```