

Base Users Guide



7.0 Edition



Base Users Guide:

7.0 Edition

Published Dec 01 2017

Copyright © 2017 University of California

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

Table of Contents

Preface	vii
1. Overview	1
2. Attributes	2
3. Installing a Rocks Cluster	5
3.1. Getting Started.....	5
3.2. Rocks 7: Install and Configure Your Frontend.....	6
3.3. Rocks 6: Install and Configure Your Frontend.....	20
3.4. Install Your Compute Nodes	34
3.5. Upgrade or Reconfigure Your Existing Frontend.....	37
3.6. Installing a Frontend over the Network.....	41
3.7. Enabling Public Web Access to Your Frontend	43
4. Defining and Modifying Networks and Network Interfaces	46
4.1. Networks, Subnets, VLANs and Interfaces.....	46
4.2. Named Networks/Subnets	46
4.3. Host Interfaces.....	46
4.4. Virtual LANs (VLANs) and Logical VLAN Bridges	47
4.5. Configuring IPMI interfaces.....	49
4.6. Networking Configuration Examples.....	51
5. Customizing your Rocks Installation	53
5.1. Adding Packages to Compute Nodes	53
5.2. Customizing Configuration of Compute Nodes	54
5.3. Adding Applications to Compute Nodes	55
5.4. Configuring Additional Ethernet Interfaces	55
5.5. Compute Node Disk Partitioning	56
5.6. Creating a Custom Kernel RPM.....	64
5.7. Enabling RSH on Compute Nodes.....	66
5.8. Hostbased vs. Key-base SSH Authentication.....	66
5.9. Adding a New Appliance Type to the Cluster.....	67
5.10. Adding a Device Driver.....	69
5.11. Extending DNS	71
5.12. Changing the Root password.....	71
5.13. The insert-ethers command line	72
6. Community Resources	74
6.1. Access to Rocks Source Code.....	74
6.2. All Past Rocks Releases	74
6.3. Email Discussion List.....	74
7. Administration Examples	76
7.1. Introduction to the Rocks Command Line	76
7.2. Boot Order and PXE First	76
7.3. Support for PXE First.....	77
7.4. Forcing a Re-install at Next PXE Boot	77
7.5. Inspecting and Changing PXE Behaviour.....	77
7.6. Working with and Modifying Network Configuration.....	79
7.7. Reinstall All Compute Nodes with SGE.....	79

8. Advanced Tasks	81
8.1. Managing the Firewall on the Cluster Nodes.....	81
8.2. Enabling Two-Factor Authentication with Google Authenticator.....	85
8.3. Flashing BIOS on Compute Nodes Using PXE.....	92
8.4. Adding a Login Appliance to Your Cluster.....	93
8.5. Channel Bonding Interfaces.....	94
8.6. Frontend Central Server.....	95
8.7. Cross Kickstarting.....	95
8.8. Adding Kernel Boot Parameters.....	97
8.9. Controlling Bootloader Parameters.....	99
8.10. System Update.....	100
8.11. Reconfigure a Cluster.....	102
9. Command Reference	105
9.1. No Commands Installed at BUILD.....	105
A. Frequently Asked Questions	106
A.1. Installation.....	106
A.2. Configuration.....	108
A.3. System Administration.....	110
B. Release Notes	112
B.1. Release 6.2 - changes from 6.1.1.....	112
B.2. Release 5.6/6.1_sp1 - changes from 5.6/6.1.....	113
B.3. Release 5.5/6.1 - changes from 5.5/6.0.....	113
B.4. Release 5.5/6.0 - changes from 5.4.3.....	114
B.5. Release 5.4.3 - changes from 5.4.....	116
B.6. Release 5.4 - changes from 5.3.....	120
B.7. Release 5.3 - changes from 5.2.....	124
B.8. Release 5.2 - changes from 5.1.....	126
B.9. Release 5.1 - changes from 5.0.....	128
B.10. Release 4.3 - changes from 4.2.1.....	130
B.11. Release 3.2.0 - changes from 3.1.0.....	131
B.12. Release 3.2.0 - changes from 3.1.0.....	132
B.13. Release 3.1.0 - changes from 3.0.0.....	133
B.14. Release 3.0.0 - changes from 2.3.2.....	133
B.15. Release 2.3.2 - changes from 2.3.1.....	134
B.16. Release 2.3.1 - changes from 2.3.....	134
B.17. Release 2.2.1 - changes from 2.2.....	135
B.18. Release 2.2 - changes from 2.1.2.....	135
B.19. Release 2.1.2 - changes from 2.1.1.....	136
B.20. Release 2.1.1 - changes from 2.1.....	136
B.21. Release 2.1 - changes from 2.0.1.....	137
B.22. Release 2.0.1 - changes from 2.0.....	138
C. 411 Secure Information Service Internals	139
C.1. Using the 411 Service.....	139
C.2. Structure.....	139
C.3. 411 Groups.....	140
C.4. Plugins.....	141
C.5. 411get Configuration File.....	142
C.6. Commands.....	143

D. Changes to Rocks Security Infrastructure	145
D.1. Rocks Password Infrastructure	145
D.2. Rocks Secure Attribute Infrastructure	146
E. Kickstart Nodes Reference	149
E.1. Rocks Base Nodes	149
F. Rocks® Copyright and Trademark	177
F.1. Copyright Statement	177
F.2. Trademark Licensing	178
G. Common Licenses	179
G.1. Artistic License	179
G.2. Apache v2.0	181
G.3. GNU General Public License v1	184
G.4. GNU General Public License v2	189
G.5. GNU Lesser General Public License v2.1	195
G.6. GNU Library General Public License v2	204
G.7. Python Software Foundation License v2	212
H. Package Licenses	214
H.1. anaconda	214
H.2. ant	214
H.3. coreutils	216
H.4. cvs	216
H.5. eGenix mx	217
H.6. FireFox	218
H.7. gawk	228
H.8. gd	228
H.9. graphviz	229
H.10. kudzu	233
H.11. libxml2	233
H.12. libxml2doc	234
H.13. mysql	234
H.14. ncurses	237
H.15. numarray	237
H.16. Numeric	238
H.17. perl	239
H.18. perl tk	240
H.19. pexpect	242
H.20. phpMyAdmin	242
H.21. POW	242
H.22. pygtk	243
H.23. python	243
H.24. rcs	243
H.25. readline	243
H.26. tidy	244
H.27. wget	244

List of Tables

1-1. Summary.....	1
1-2. Compatibility.....	1
2-1. Roll Attributes	2
3-1. Frontend -- Default Root Disk Partition	31
5-1. Compute Node -- Default Root Disk Partition	56
5-2. A Compute Node with 3 SCSI Drives	62

Preface

Since May 2000, the Rocks group has been addressing the difficulties of deploying manageable clusters. We have been driven by one goal: *make clusters easy*. By *easy* we mean easy to deploy, manage, upgrade and scale. We are driven by this goal to help deliver the computational power of clusters to a wide range of scientific users. It is clear that making stable and manageable parallel computing platforms available to a wide range of scientists will aid immensely in improving the state of the art in parallel tools.

Chapter 1. Overview

Table 1-1. Summary

Name	base
Version	7.0
Maintained By	Rocks Group
Architecture	x86_64
Compatible with Rocks®	7.0

The base roll has the following requirements of other rolls. Compatability with all known rolls is assured, and all known conflicts are listed. There is no assurance of compatiblity with third-party rolls.

Table 1-2. Compatibility

Requires	Conflicts
Base	
Kernel	

Chapter 2. Attributes

Table 2-1. Roll Attributes

Name	Type	Default
disableServices	string	kudzu canna cWnn FreeWnn kWnn tWnn mDNSResponder
Info_CertificateCountry a	string	
Info_CertificateLocality a	string	
Info_CertificateOrganization a	string	
Info_CertificateState a	string	
Info_CertificateContact a	string	
Info_CertificateLatLong a	string	
Info_CertificateName a	string	
Info_CertificateURL a	string	
Kickstart_DistroDir a	string	/export/rocks
Kickstart_Keyboard a	string	us
Kickstart_Lang a	string	en_US
Kickstart_Langsupport a	string	en_US
Kickstart_Mutlicast a	string	226.117.172.185
Kickstart_PrivateAddress a	string	10.1.1.1
Kickstart_PrivateBroadcast a	string	10.1.255.255
Kickstart_PrivateDNSDomain a	string	local
Kickstart_PrivateDNSServers a	string	10.1.1.1
Kickstart_PrivateGateway a	string	10.1.1.1
Kickstart_PrivateHostname a	string	
Kickstart_PrivateKickstartBaseDir a	string	install
Kickstart_PrivateKickstartCGI a	string	sbin/kickstart.cgi
Kickstart_PrivateKickstartHost a	string	10.1.1.1
Kickstart_PrivateNTPHost a	string	10.1.1.1
Kickstart_PrivateNetmask a	string	255.255.0.0
Kickstart_PrivateNetmaskCIDR a	string	16
Kickstart_PrivateNetwork a	string	10.1.0.0
Kickstart_PrivatePortableRootPassword a	string	
Kickstart_PrivateRootPassword a	string	
Kickstart_PrivateSHARootPassword a	string	
Kickstart_PrivateSyslogHost a	string	10.1.1.1
Kickstart_PublicAddress a	string	
Kickstart_PublicBroadcast a	string	
Kickstart_PublicDNSDomain a	string	
Kickstart_PublicDNSServers a	string	
Kickstart_PublicGateway a	string	
Kickstart_PublicHostname a	string	

Name	Type	Default
Kickstart_PublicKickstartHost ^a	string	
Kickstart_PublicNTPHost ^a	string	
Kickstart_PublicNetmask ^a	string	
Kickstart_PublicNetmaskCIDR ^a	string	
Kickstart_PublicNetwork ^a	string	
Kickstart_Timezone ^a	string	
airboss ^b	string	specified on boot line
arch ^{c, b}	string	i386 x86_64
dhcp_filename ^d	string	pxelinux.0
dhcp_nextserver ^d	string	10.1.1.1
hostname ^{e, b}	string	
kickstartable ^d	bool	TRUE
os ^{c, b}	string	linux solaris
rack ^{e, b}	int	
rank ^{e, b}	int	
rocks_version ^a	string	7.0
rsh ^f	bool	FALSE
rocks_autogen_user_keys ^f	bool	FALSE
ssh_use_dns ^a	bool	TRUE
x11 ^f	bool	FALSE

Notes:

- Default value created using **rocks add attr name value** and affects all hosts.
- Default value created using **rocks add host attr localhost name value** and only affects the `frontend` appliance.
- Attribute is for internal use only, and should not be altered by the user. Each time a machine installs this attributed is reset to the default value for that machine (depend on kernel booted).
- Default value created using **rocks add appliance attr appliance name value** for the `frontend` and `compute` appliances.
- Attribute cannot be modified. This value is not recorded in the cluster database and is only available as an XML entity during installation.
- Attribute is referenced but not defined so is treated as `FALSE`.

Info_Certificate_{*}

The attributes are created during `frontend` installation. The values are taken from user input on the system installation screens.

Kickstart_{*}

The attributes are created during `frontend` installation. The values are taken from user input on the system installation screens. All of these attributes are considered internal to Rocks® and should not be modified directly.

airboss

Specifies the address of the airboss host. This only applies to virtual machines.

`arch`

The CPU architecture of the host. This host-specific attribute is set by the installing machine. User changes to this attribute have no affect.

`dhcp_filename`

Name of the PXE file retrieved over TFTP at startup.

`dhcp_nextserver`

IP address of the server that servers installation profiles (kickstart, jumpstart). In almost all configuration this should be the `frontend` machine.

`kickstartable`

The attribute must be set to `TRUE` for all appliances, and `FALSE` (or undefined) for all unmanaged devices (e.g. network switches).

`os`

The OS of the host. This host-specific attribute is set by the installing machine. User changes to this attribute have no affect.

`rsh`

If `TRUE` the machine is configured as an RSH client. This is not recommended, and will still require RSH server configuration on the `frontend` machine.

`ssh_use_dns`

Set to `FALSE` to disable DNS lookups when connecting to nodes in the cluster over SSH. If establishing an ssh connect is slow the cause may be a faulty (or absent) DNS system. Disabling this lookup will speed up connection establishment, but lowers the security of your system.

`x11`

If `TRUE` X11 is configured and the default runlevel is changed from 3 to 5. X11 is always configure on the `frontend` and this attribute applies only to the other nodes in the cluster.

Chapter 3. Installing a Rocks Cluster

3.1. Getting Started

This chapter describes the steps to build your cluster and install its software.

3.1.1. Supported Hardware

Processors

- x86 (ia32, AMD Athlon, etc.)
- x86_64 (AMD Opteron and EM64T)

Networks

- Ethernet



Specialized networks and components (e.g., Myrinet, Infiniband, nVidia GPU) are also supported. Hardware requirements and software (Rocks Rolls) can be found on the respective vendor web sites.

3.1.2. Minimum Hardware Requirements

Frontend Node

- **Disk Capacity:** 36 GB
- **Memory Capacity:** 2 GB
- **Ethernet:** 2 physical ports (e.g., "eth0" and "eth1")
- **BIOS Boot Order:** CD, Hard Disk

Compute Node


- **Disk Capacity:** 36 GB
- **Memory Capacity:** 2 GB
- **Ethernet:** 1 physical port (e.g., "eth0")
- **BIOS Boot Order:** PXE (Network Boot), Hard Disk

3.1.3. Physical Assembly

The first thing to manage is the physical deployment of a cluster. Much research exists on the topic of how to physically construct a cluster. A majority of the O'Reilly Book¹ *Building Linux Clusters* is devoted to the physical setup of a cluster, how to choose a motherboard, etc. Finally, the book *How to Build a Beowulf* also has some good tips on physical construction.

We favor rack-mounted equipment because of its relative reliability and density. There are Rocks clusters, however, that are built from mini-towers. Choose what makes sense for you.

The following diagram shows how the frontend and compute nodes must be connected:



On the compute nodes, the Ethernet interface that Linux maps to `eth0` should be connected to the cluster's Ethernet switch. This network is considered *private*, that is, all traffic on this network is physically separated from the external public network (e.g., the internet).

On the frontend, at least two ethernet interfaces are required. The interface that Linux maps to `eth0` should be connected to the same ethernet network as the compute nodes. The interface that Linux maps to `eth1` should be connected to the external network (e.g., the internet or your organization's intranet).

3.2. Rocks 7: Install and Configure Your Frontend

This section describes how to install your Rocks cluster frontend for Rocks Version 7. It is significantly different than Rocks Version 6.



The minimum requirement to bring up a frontend is to have the following rolls:

- Kernel Roll
- Base Roll
- Core Roll
- CentOS Roll
- Updates-CentOS Roll

Rocks 7 has only been tested on CentOS. Other versions, like RedHat at the same version number may also work.



Rocks version 7 supports a network-only installation. All rolls must be located on a roll server on a network that is accessible by your frontend

1. Insert the Kernel Roll CD into your frontend machine and reset the frontend machine.

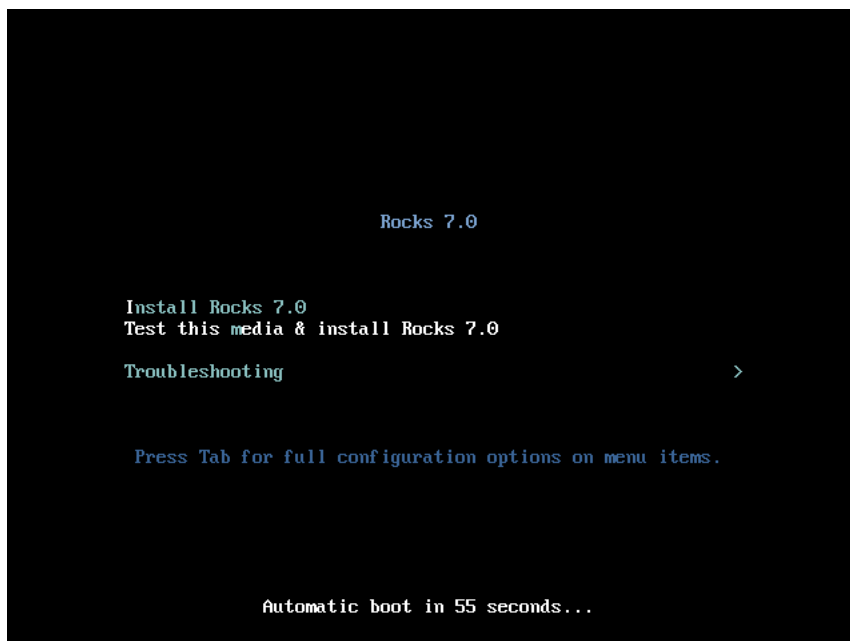


You may install from a USB stick. Simply copy the kernel iso image to a usb stick using `dd`. It is compatible with both UEFI and legacy BIOS systems. The standard process of burning a CentOS USB bootable stick² from the kernel iso image



For the remainder of this section, we'll use the example of installing a fairly complete frontend, with multiple rolls, that is, we'll be using

2. After the frontend boots off the CD or bootable usb, you will see:



When you see the screen above, you can select "Install Rocks 7.0" or "Test this media & install Rocks 7.0". Furthermore, you can add boot options. The Rocks 7.0 installer is built using extensions to anaconda, RedHat/CentOS installer and anaconda boot options are available



Unlike, Rocks 6, this bootable CD is only used for frontends. Client nodes must install over the network using PXE



All of the anaconda installer terminals are available. In particular, `Ctrl-Alt-F2` accesses a shell prompt on your system respectively.

3. Select language and configure your public network

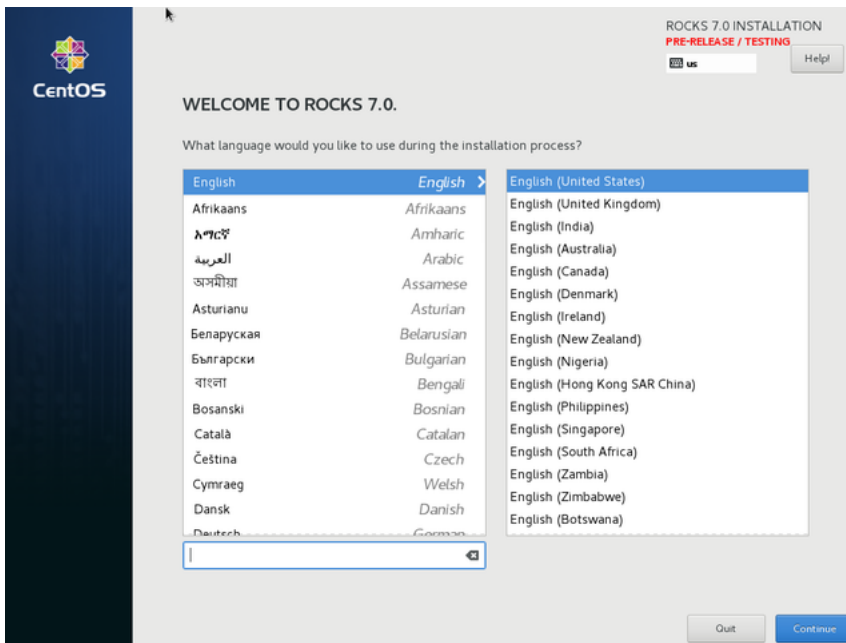


The CentOS installer, does not require a specific order to the screens of information that must be entered. If a certain selection is grayed-out, it indicates that the particular section is not ready to be visited



Not all sections are visible you must use the *scrollbar* to see them all. In particular, it is quite common to forget to properly set the timezone

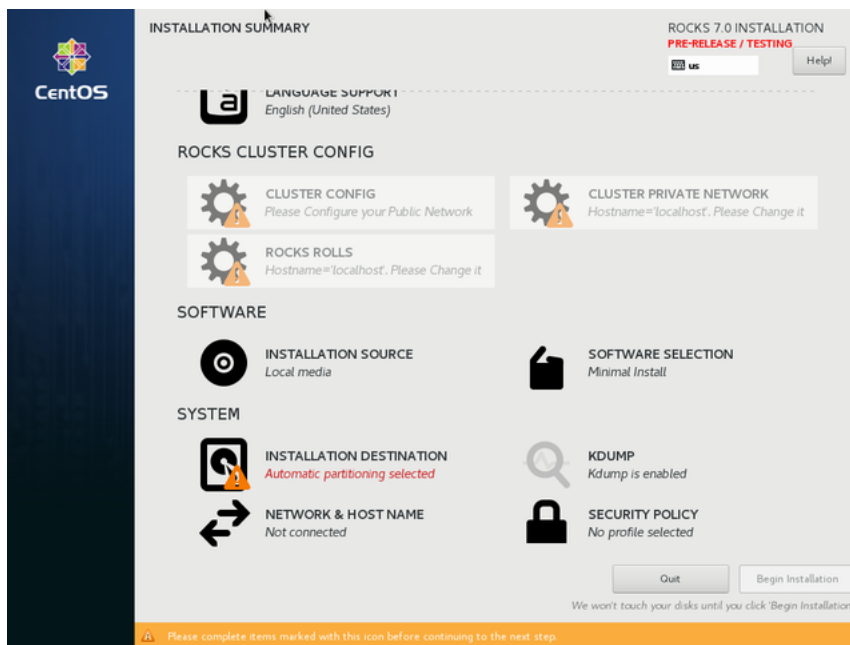
Your first see, the language screen. Rocks 7.0 has only been tested with English/US. Other language selections may cause issues



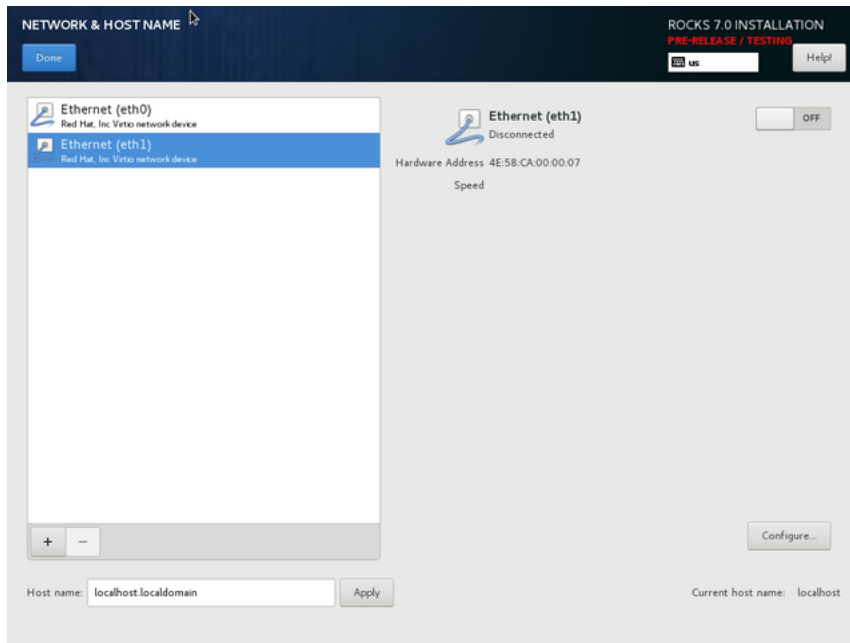
If your local network does not automatically supply you with a network configuration, you will see a screen that looks like the following. The grayed selections indicate that something else must be done. In this case, the public network must be configured and the hostname must not be localhost



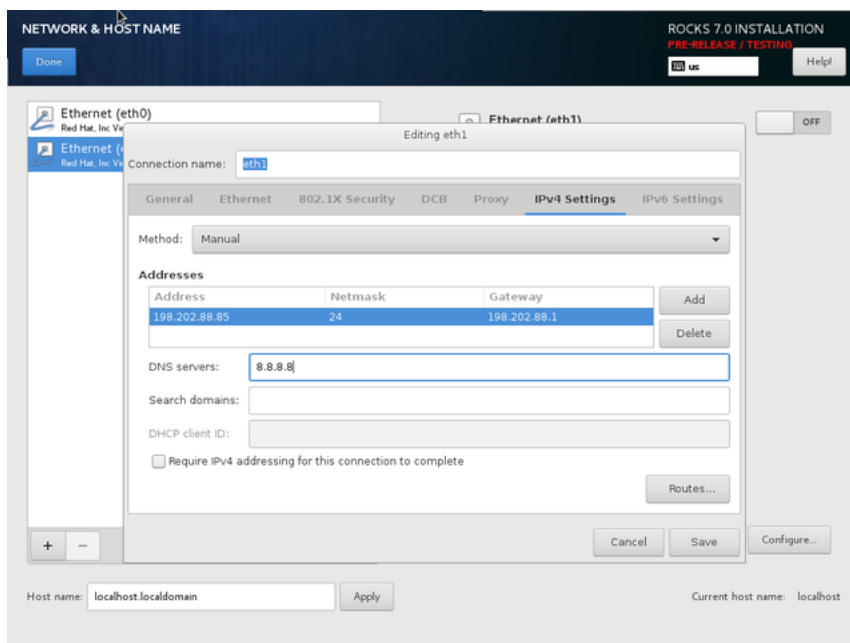
The network configuration screen is in the "System Area" and you need to scroll using the scrollbar supplied by anaconda. This scrollbar can be "thin" or "difficult to see". The following picture illustrates the scrolled down.



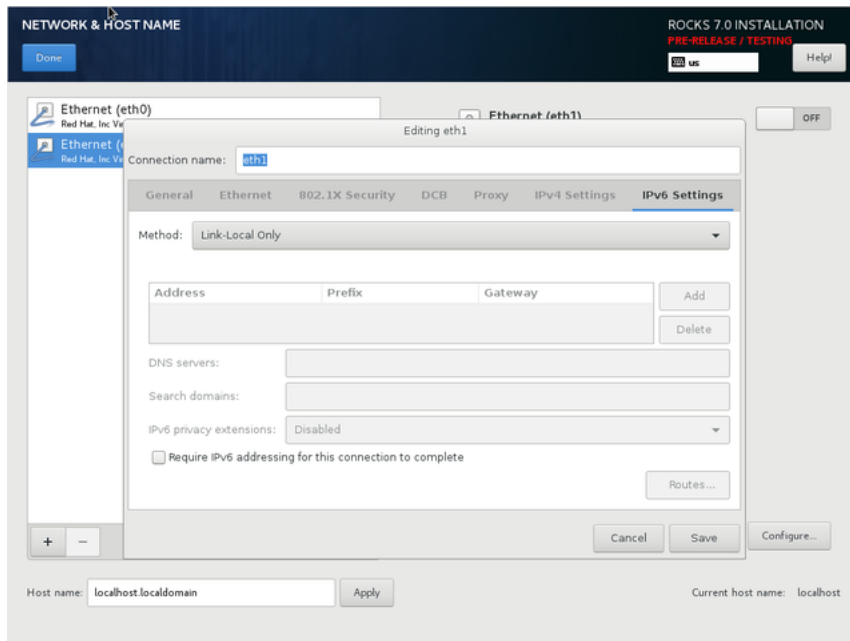
Click on the "NETWORK & HOST NAME" area to configure your public IP and hostname. If you have multiple interfaces, configure *only the interface for public access* on this screen. It is `eth1` in this example.



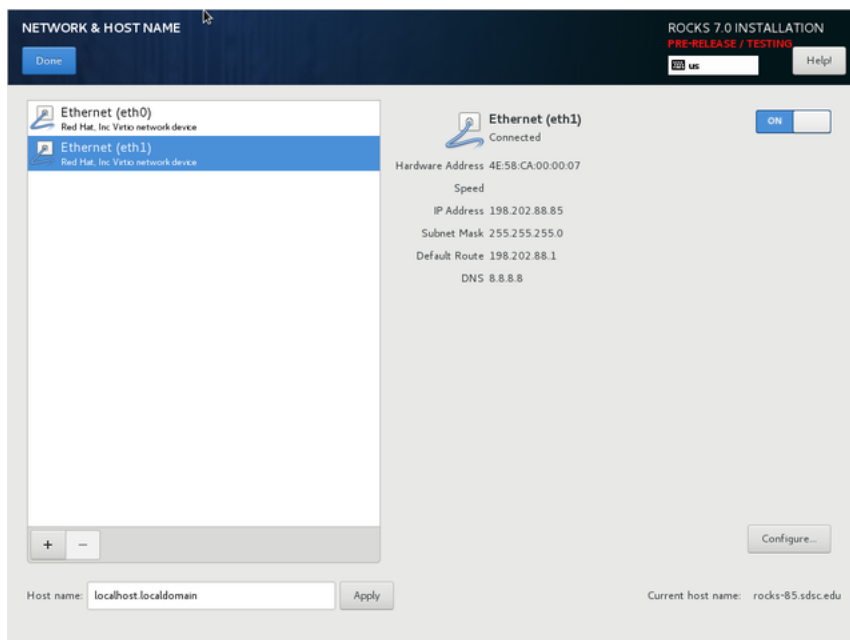
click the "Configure..." button to get to network configuration. In the following screen, configure the IPv4 address. Notice that it is manual configuration.



Also configure IPv6 to "link-local only"



After clicking "Save" and "Done" you return to the network screen. At this point, click the on/off slider in the upper right-hand corner of the screen to activate your network interface. If all goes well, you should see both your IP address and DNS name. If you desire to change the hostname, edit the Host name box and click "Apply"

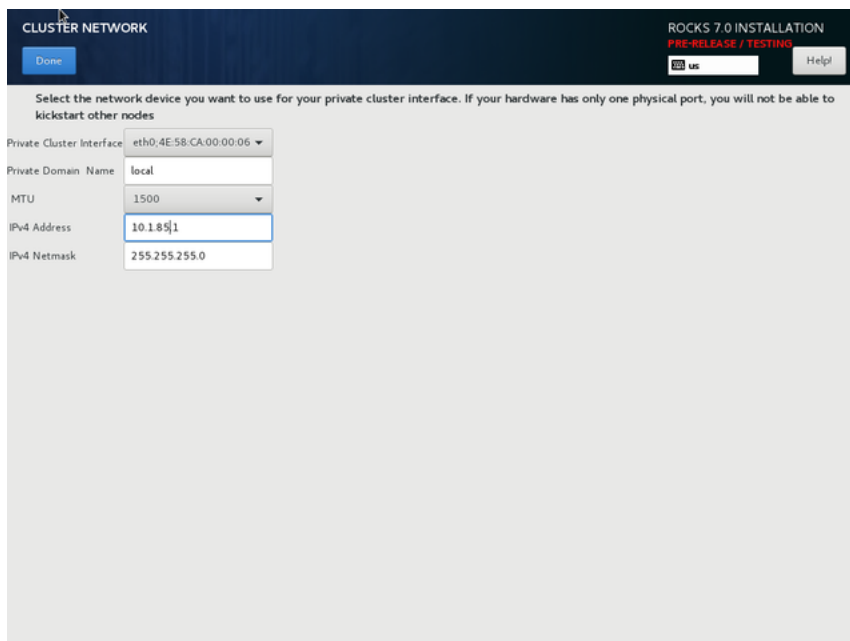


If you are successful, your home screen should now reflect that the network is up and configured.



4. Configure your local network.

Scroll up to the ROCKS CLUSTER CONFIG area and then click on the CLUSTER PRIVATE NETWORK screen.



The "private cluster interface" is a drop-down menu selection that allows you to choose any physical interface for the private network. The interface you chose for your public network is not available on this list. If your system has only one physical interface, a virtual interface is constructed for you.



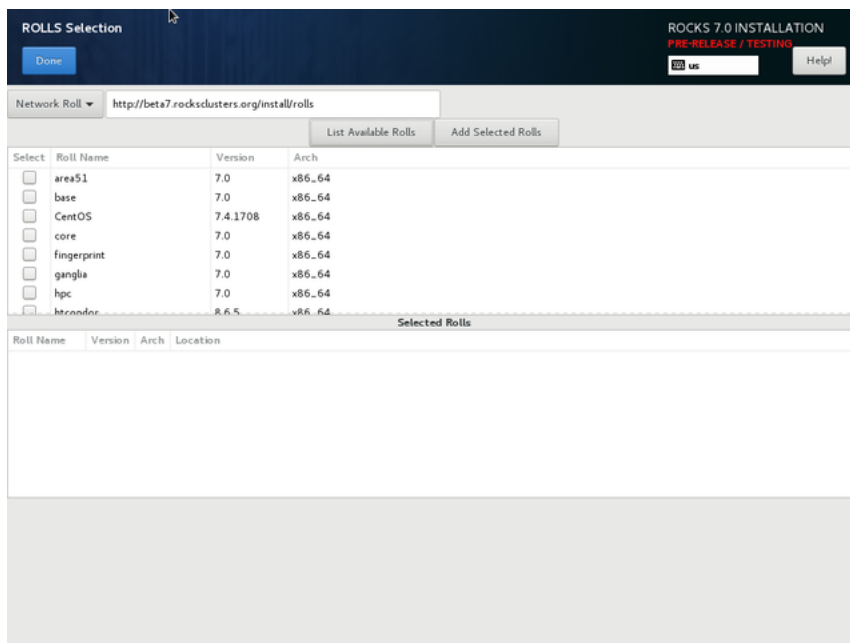
You may select any IPv4 subnet you choose. However, conflict with the public network is not checked. Your public and private IP subnetworks must not overlap.

When you click save, the home screen now indicates that this configuration section is complete.



5. Select your rolls. Click on the ROCKS ROLLS selection

Rolls define your cluster’s configuration. The default location for rolls may be suboptimal for your cluster in terms of accessibility or performance. It is possible to mirror the rolls onto any (non-rocks) http (web) server. When you have selected your roll server, click the `List Available Rolls` button.



You have several choices in selecting rolls. You may hit the "Select" word at the top of the column of check-boxes. This will select all rolls available at that server. You may select/de-select individual rolls, too.

ROCKS 7.0 INSTALLATION
PRE-RELEASE / TESTING

Network Roll: <http://beta7.rocksclusters.org/install/rolls>

Select	Roll Name	Version	Arch
<input checked="" type="checkbox"/>	area51	7.0	x86_64
<input checked="" type="checkbox"/>	base	7.0	x86_64
<input checked="" type="checkbox"/>	CentOS	7.4.1708	x86_64
<input checked="" type="checkbox"/>	core	7.0	x86_64
<input checked="" type="checkbox"/>	fingerprint	7.0	x86_64
<input checked="" type="checkbox"/>	ganglia	7.0	x86_64
<input checked="" type="checkbox"/>	hpc	7.0	x86_64
<input checked="" type="checkbox"/>	htcondor	8.6.5	x86_64

Selected Rolls

Roll Name	Version	Arch	Location
-----------	---------	------	----------



You must click on "Add Selected Rolls" to actually add these rolls for your frontend

ROCKS 7.0 INSTALLATION
PRE-RELEASE / TESTING

Network Roll: <http://beta7.rocksclusters.org/install/rolls>

Select	Roll Name	Version	Arch
--------	-----------	---------	------

Selected Rolls

Roll Name	Version	Arch	Location
area51	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls
base	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls
CentOS	7.4.1708	x86_64	http://beta7.rocksclusters.org/install/rolls
core	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls
fingerprint	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls
ganglia	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls
hpc	7.0	x86_64	http://beta7.rocksclusters.org/install/rolls



You can download rolls from multiple roll servers. Simple edit the URL of another roll server and click "List Available Rolls".

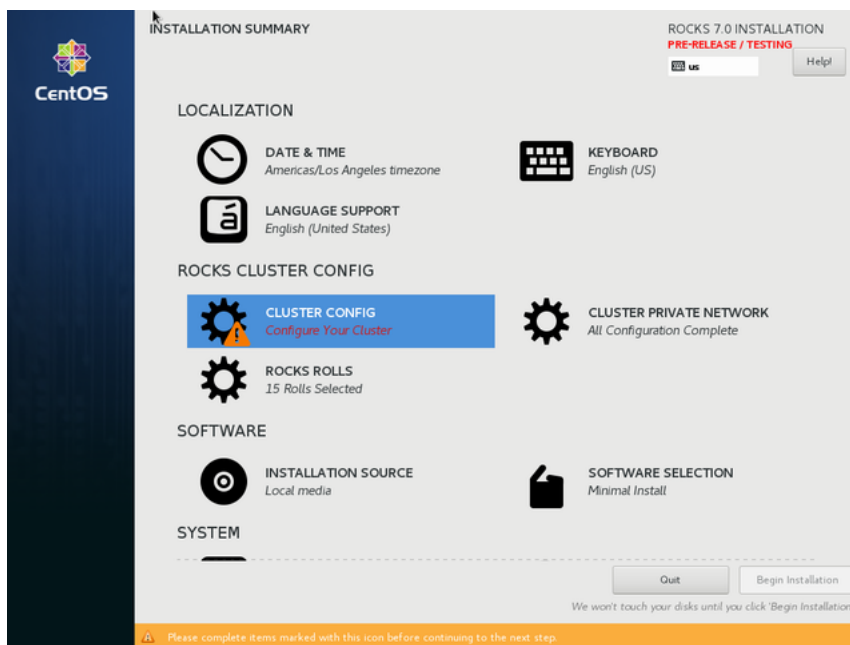
When you hit done and you have selected the required base, core, and kernel rolls, it will take 5-10 seconds for Rocks to build a local database. At this point the screen will look like.



This indicates that 15 different rolls have been selected.

6. Configure your Cluster

You must enter the CLUSTER CONFIG screen to finish configuration



Any items that need to be completed are highlighted in red.

Cluster Config ROCKS 7.0 INSTALLATION
PRE-RELEASE / TESTING
US Help

Done

Configuration Information. You must complete all required entries marked (in red, if invalid). Grey colored entries are entered elsewhere in the installer and shown here for brevity. Hover over each line to get a longer detailed explanation

parameter	value
Fully-Qualified Domain Name	rocks-85.sdsc.edu
Cluster Name	rocks85
Contact	admin@place.org
Project URL	http://www.place.org
Latitude/Longitude	N32.86 W117.22
Certificate Organization	SDSC
Certificate Locality	San Diego
Certificate State	California
Certificate Country	US
Network Device for Private Network	eth0
Private Cluster IPv4 Address	10.1.85.1
Private Cluster Netmask	255.255.255.0
Private Cluster Interface MTU	1500
Timezone	America/Los_Angeles
NTP Servers	pool.ntp.org
Public Interface	eth1
Public IPv4 Netmask	255.255.255.0
Public Interface MTU	1500
Default Router for Public IPv4 Network	198.202.88.1

Once a contact has been filled in, the red will disappear. Please also note that grey areas are informational and is data gathered from other parts of the installer.

Cluster Config ROCKS 7.0 INSTALLATION
PRE-RELEASE / TESTING
US Help

Done

Configuration Information. You must complete all required entries marked (in red, if invalid). Grey colored entries are entered elsewhere in the installer and shown here for brevity. Hover over each line to get a longer detailed explanation

parameter	value
Fully-Qualified Domain Name	rocks-85.sdsc.edu
Cluster Name	rocks85
Contact	admin@place.org
Project URL	http://www.place.org
Latitude/Longitude	N32.86 W117.22
Certificate Organization	SDSC
Certificate Locality	San Diego
Certificate State	California
Certificate Country	US
Network Device for Private Network	eth0
Private Cluster IPv4 Address	10.1.85.1
Private Cluster Netmask	255.255.255.0
Private Cluster Interface MTU	1500
Timezone	America/Los_Angeles
NTP Servers	pool.ntp.org
Public Interface	eth1
Public IPv4 Netmask	255.255.255.0
Public Interface MTU	1500
Default Router for Public IPv4 Network	198.202.88.1

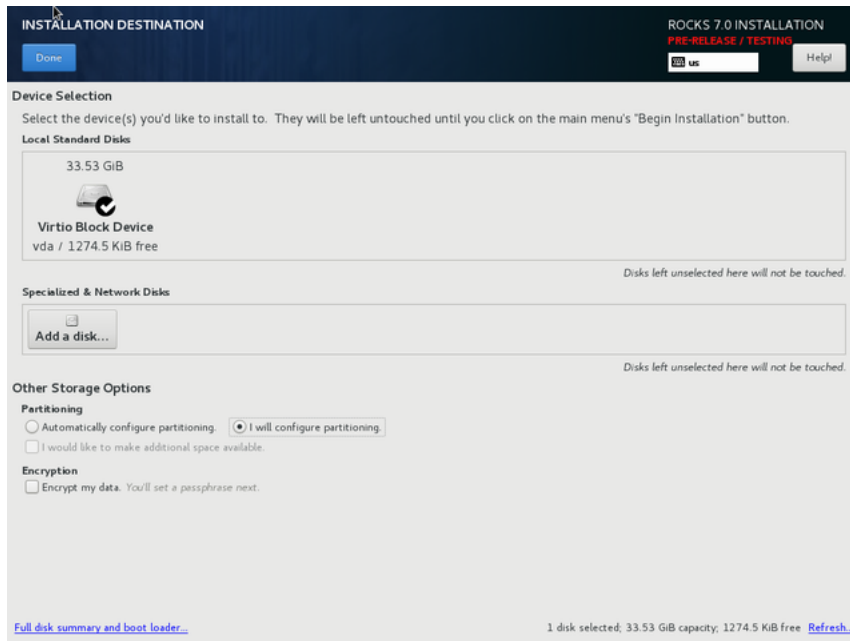
7. Configure partitioning on your frontend

Rocks 7 only supports manual partitioning on frontends. You will need to scroll down to the INSTALLATION DESTINATION screen. It will default to Automatic partitioning (which uses LVM).



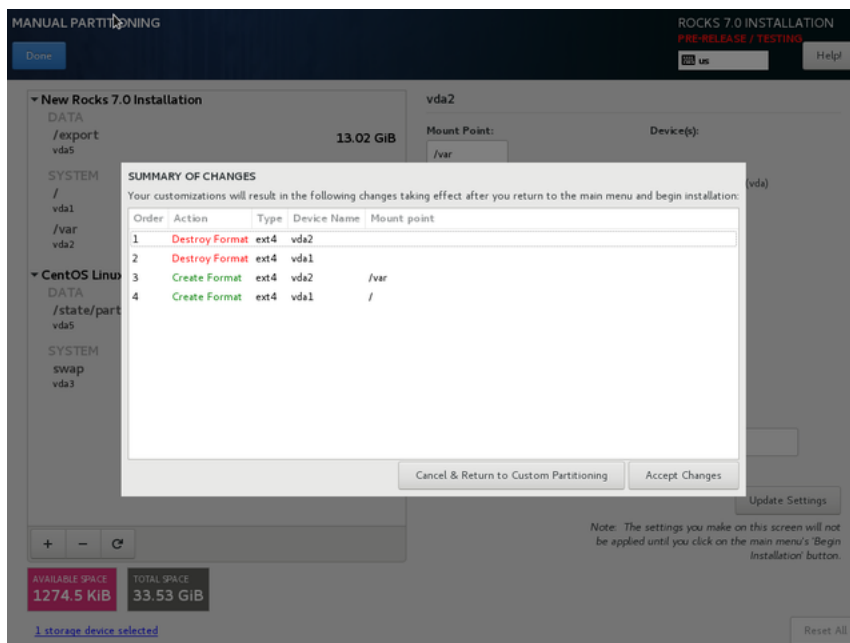
Rocks has no particular partitioning requirements on frontends, however the partition that holds your `/export/rocks` directory must be large enough to mirror all of CentOS and rolls (at least 10GB)

Highlight the "I will configure partitioning" button and then click done.



At this point, you are in Anaconda's partitioning system and can follow its usual setup. You can setup raid (recommended), or simply create one large "/" partition (not advised, but supported). It is beyond the scope of this document to highlight all of the possible disk configuration options.

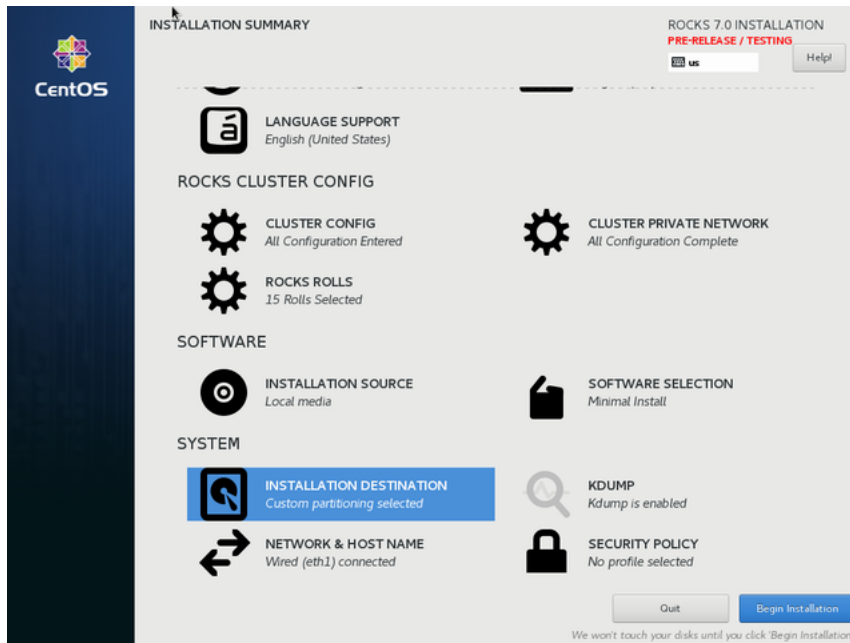
When you click "Done", you should get a confirmation screen. The following is a sample configuration.



8. Begin Installation



Once you click "Begin Installation", There is no going back. Up until this point, all changes are held in memory. After this, partitions are formatted

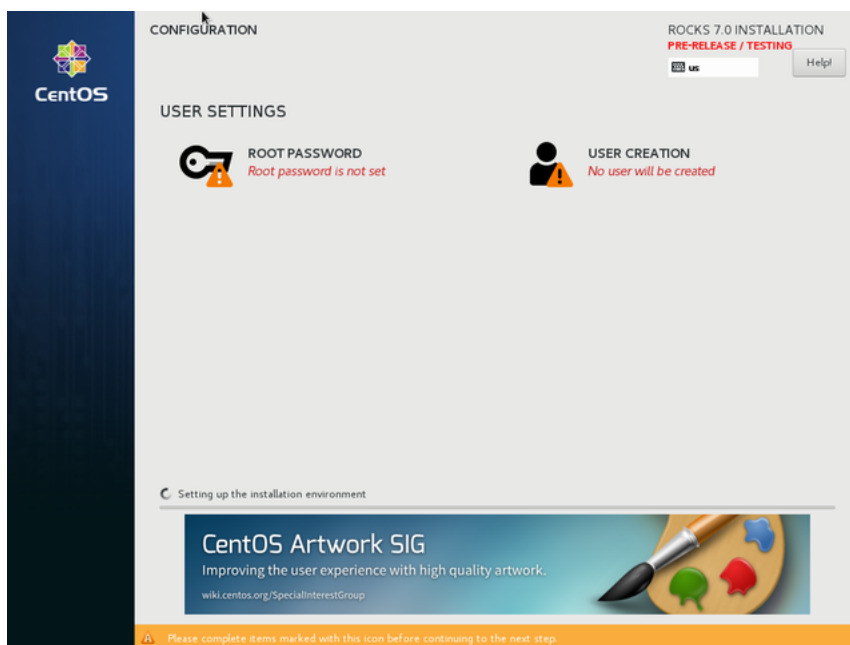


There is no "confirm" to "Begin Installation"

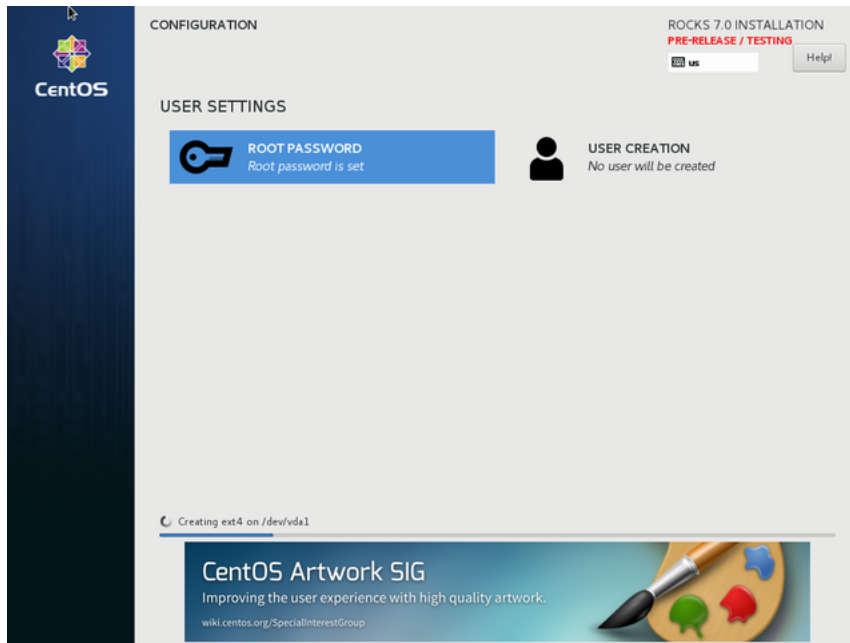
9. Let installation progress. And setup root password

Installation is now automated, but there is a lot of work to do. Partitions must be formatted, rolls must be downloaded, packages are installed and then your system is configured

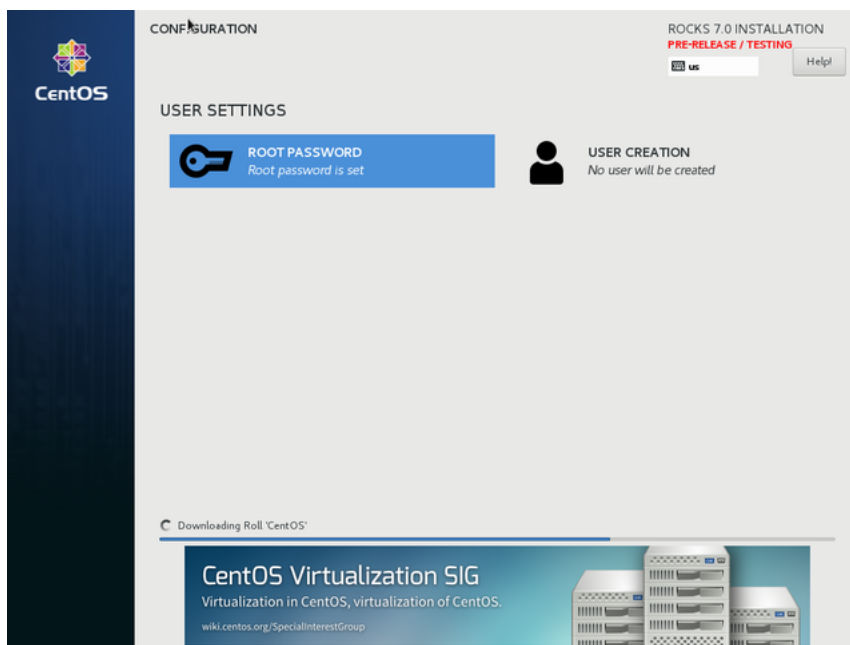
Click on the SET ROOT PASSWORD item and set your password. Do NOT setup a user at this point



Your installation should now look similar to



The above shows an ext4 file system being created. After that rolls are downloaded

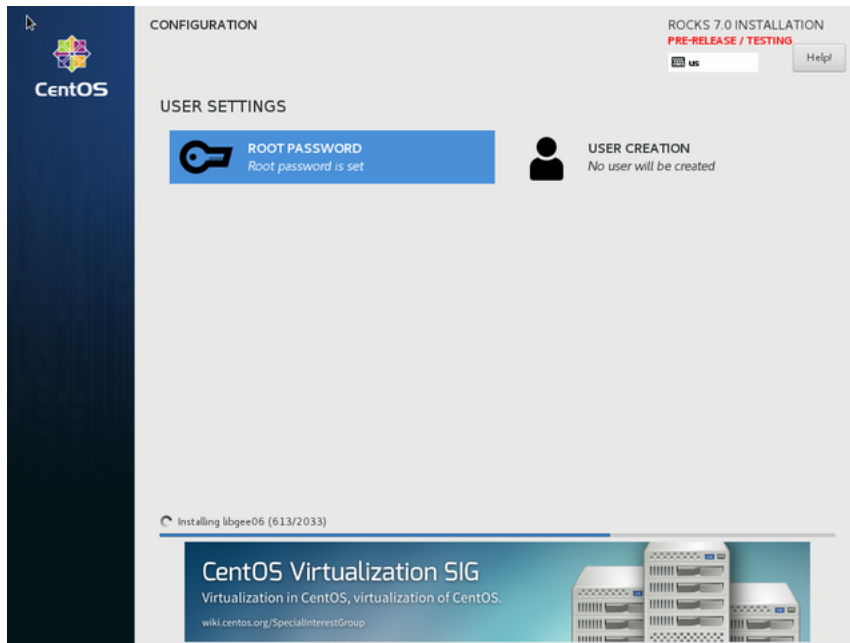


It will take a while to download all the rolls you selected, depending on your network speed. You must be patient. If you bring up the command console (CTRL-ALT-F2), you can follow the download progress with

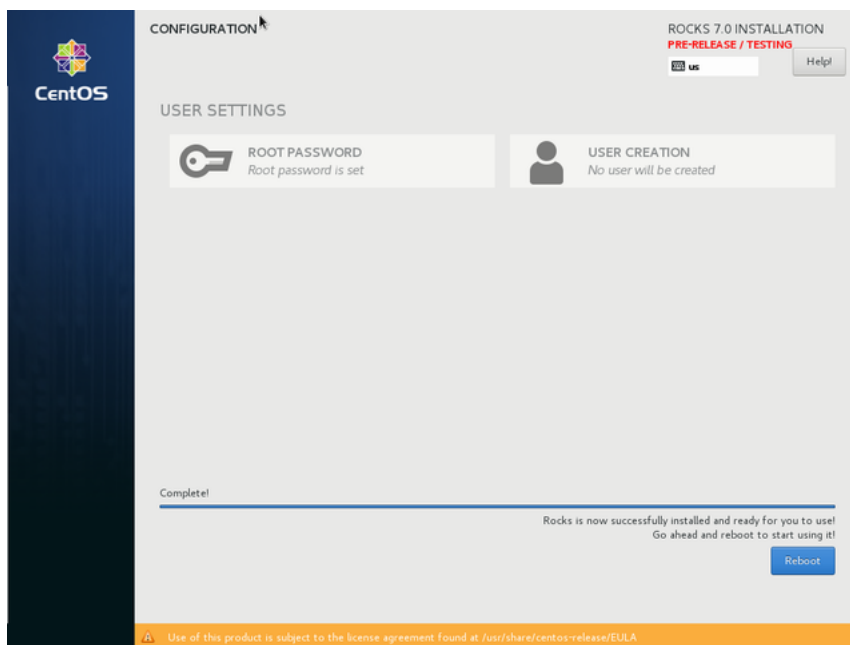
```
# find /mnt/sysimage/export/rocks/install/rolls | wc
```

The CentOS roll is over 9000 packages

After a while, all rolls are downloaded, the local distribution is created, and packages are installing



After post-installation is complete, you will see the final screen



Click reboot, remove your bootable CD or USB stick, and then wait for your system to come up

3.3. Rocks 6: Install and Configure Your Frontend

This section describes how to install your Rocks cluster frontend.



The minimum requirement to bring up a frontend is to have the following rolls:

- Kernel/Boot Roll CD

- Base Roll CD
- OS Roll CD - Disk 1
- OS Roll CD - Disk 2

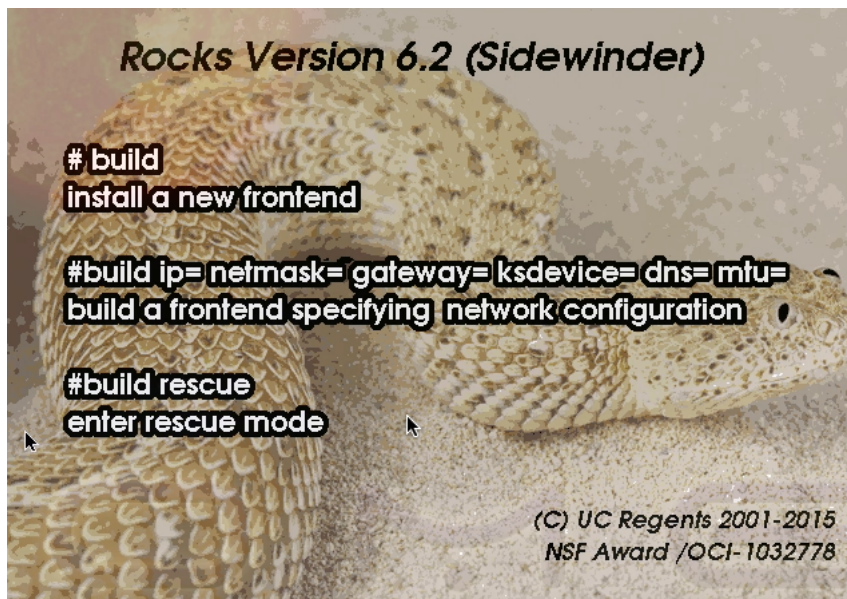
Additionally, the official Red Hat Enterprise Linux 5 (5.8) or Linux (6.3) can be substituted substituted for the OS Rolls. Also, any *true* rebuild of RHEL 5 update 8 or RHEL 6 Update 3 can be used. If you substitute the OS Rolls with one of the above distributions, you must supply *all* the CDs from the distribution (which usually is 6 to 9 CDs).

1. Insert the Kernel/Boot Roll CD into your frontend machine and reset the frontend machine.



For the remainder of this section, we'll use the example of installing a *bare-bones* frontend, that is, we'll be using the Kernel/Boot Roll, base Roll, OS - Disk 1 Roll and the OS - Disk 2 Roll.

2. After the frontend boots off the CD, you will see:



When you see the screen above, type:

```
build
```



The “boot:” prompt arrives and departs the screen quickly. It is easy to miss. If you do miss it, the node will assume it is a *compute* appliance, and the frontend installation will fail and you will have to restart the installation (by rebooting the node).



It is possible to bypass the DHCP process and have the install ask for network parameters. If you know the name of device used by the kernel for public access (e.g. `eth1,p2p1,...`), then specify as follows (using `p2p1` for the public net): `build ksdevice=p2p1 asknetwork` respectively.



If the installation fails, very often you will see a screen that complains of a missing `/tmp/ks.cfg` kickstart file. To get more information about the failure, access the kickstart and system log by pressing `Ctrl-Alt-F3` and `Ctrl-Alt-F4` respectively.

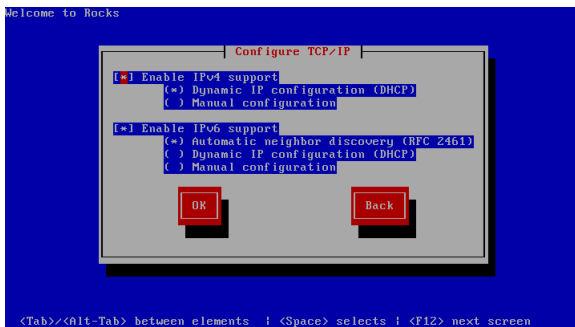
After you type `build`, the installer will start running.

3.

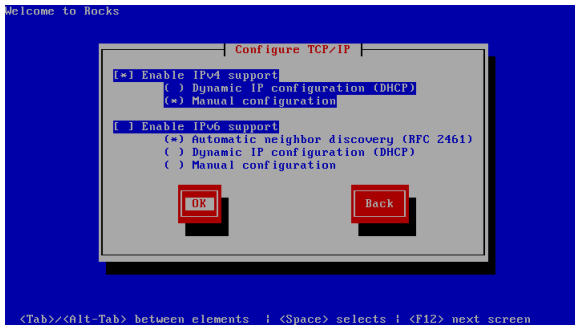


All screens in this step may not appear during your installation. You will only see these screens if there is not a DHCP server on your public network that answers the frontend's DHCP request.

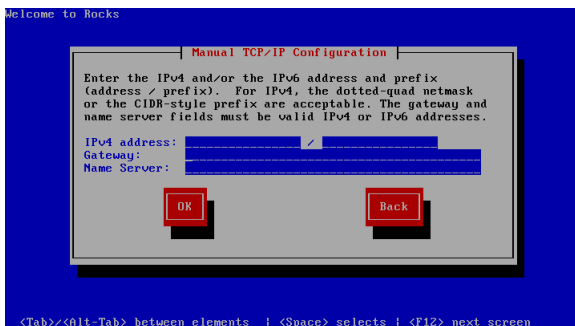
If you see the screen below:




You'll want to: 1) enable IPv4 support, 2) select manual configuration for the IPv4 support (no DHCP) and, 3) disable IPv6 support. The screen should look like:



After your screen looks like the above, hit "OK". Then you'll see the "Manual TCP/IP Configuration" screen:

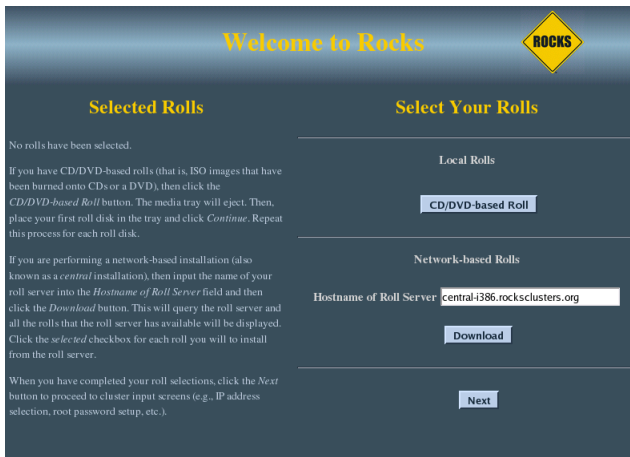


In this screen, enter the *public* IP configuration. Here's an example of the public IP info we entered for one our frontends:



After you fill in the public IP info, hit "OK".

4. Soon, you'll see a screen that looks like:

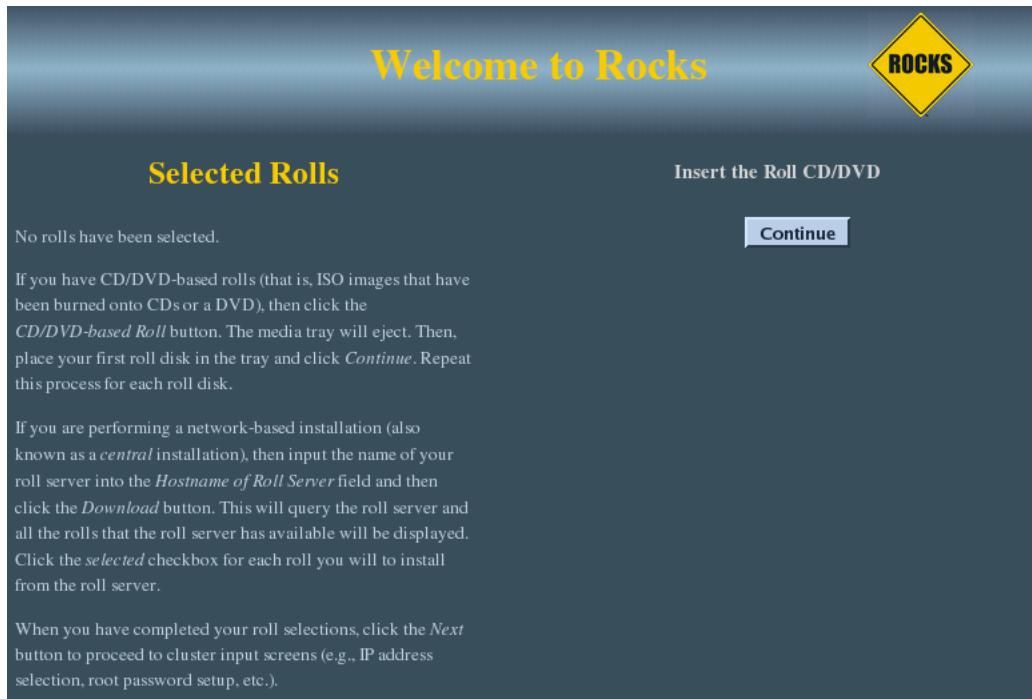


From this screen, you'll select your rolls.

In this procedure, we'll only be using CD media, so we'll only be clicking on the 'CD/DVD-based Roll' button.

Click the 'CD/DVD-based Roll' button.

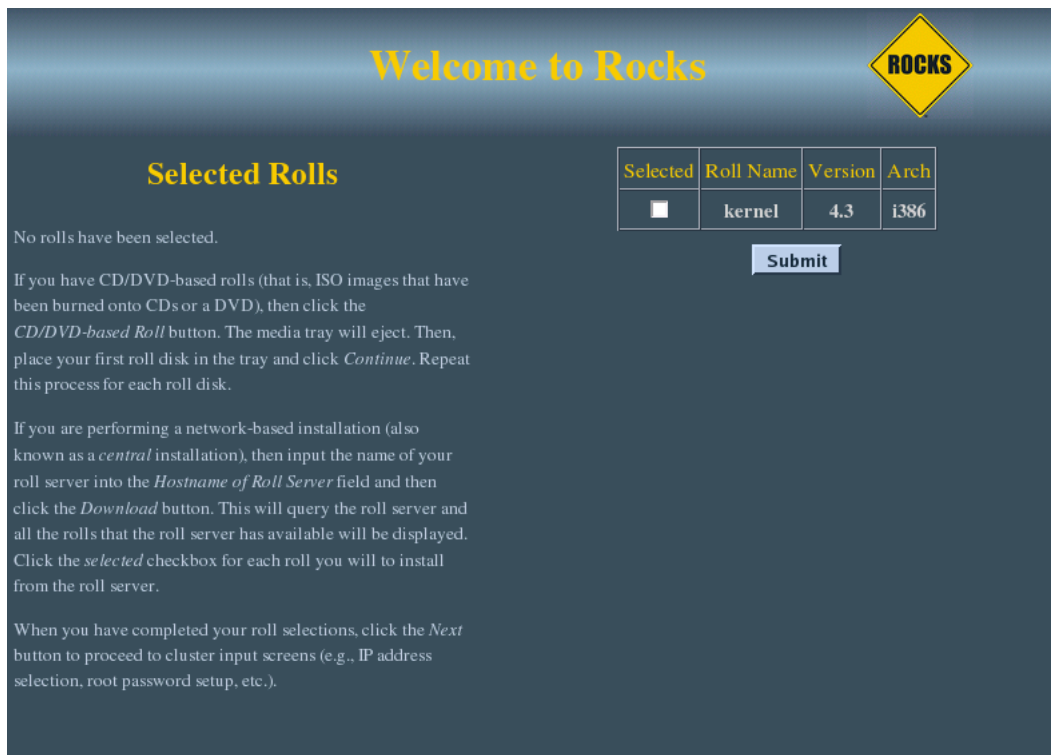
5. The CD will eject and you will see this screen:



Put your first roll in the CD tray (for the first roll, since the Kernel/Boot Roll is already in the tray, simply push the tray back in).

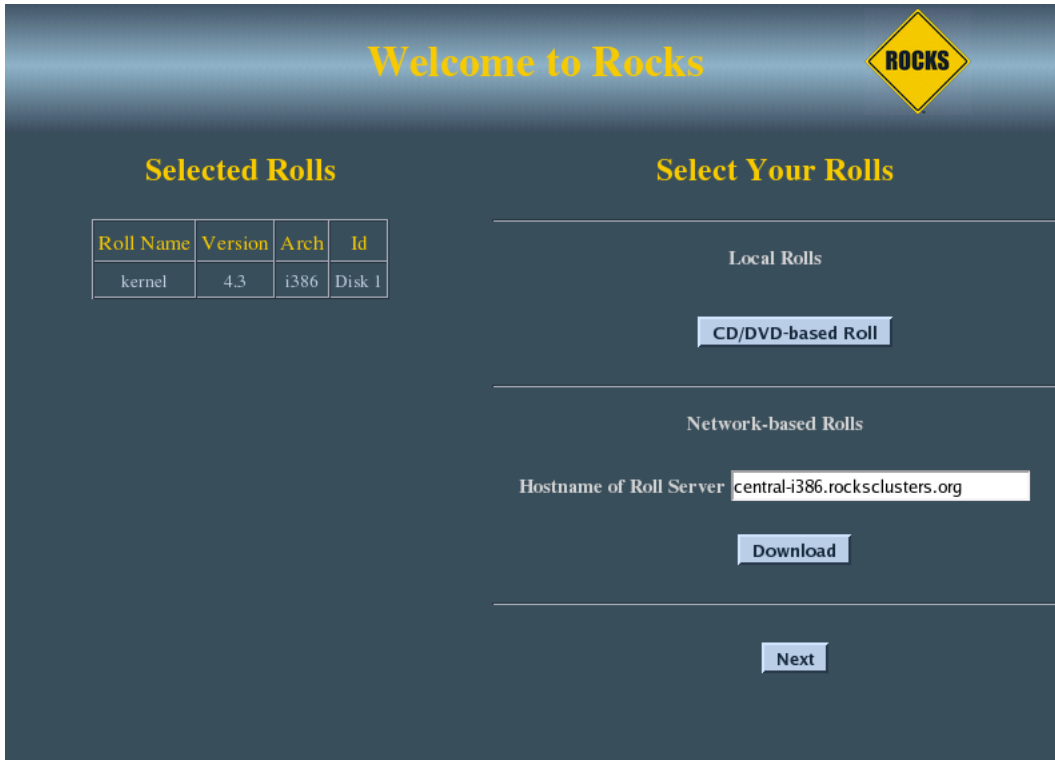
Click the 'Continue' button.

6. The Kernel/Boot Roll will be discovered and display the screen:



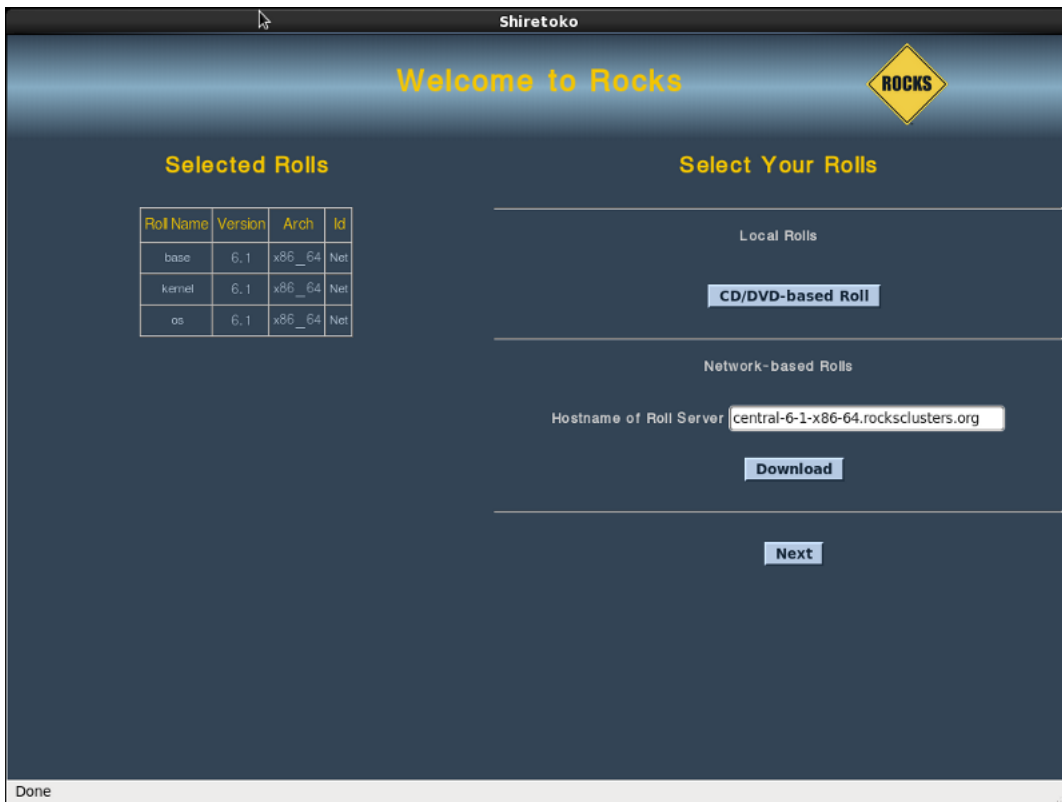
Select the Kernel/Boot Roll by checking the 'Selected' box and clicking the 'Submit' button.

7. This screen shows you have properly selected the Kernel/Boot Roll.



Repeat steps 3-5 for the Base Roll and the OS rolls.

- When you have selected all the rolls associated with a *bare-bones* frontend, the screen should look like:



When you are done with roll selection, click the 'Next' button.

- Then you'll see the *Cluster Information* screen:

Welcome to Rocks

ROCKS

Help

Fully-Qualified Host Name:
This must be the fully-qualified domain name (required).

Cluster Name:
The name of the cluster (optional).

Certificate Organization:
The name of your organization. Used when building a certificate for this host (optional).

Certificate Locality:
Your city (optional).

Certificate State:
Your state (optional).

Certificate Country:

Cluster Information

Fully-Qualified Host Name: cluster.hpc.org

Cluster Name: Our Cluster

Certificate Organization: SDSC

Certificate Locality: San Diego

Certificate State: California

Certificate Country: US

Contact: admin@place.org

URL: http://www.place.org/

Latitude/Longitude: N32.87 W117.22

Back Next



The one important field in this screen is the *Fully-Qualified Host Name* (all other fields are optional).

Choose your hostname carefully. The hostname is written to dozens of files on both the frontend and compute nodes. If the hostname is changed after the frontend is installed, several cluster services will no longer be able to find the frontend machine. Some of these services include: SGE, NFS, AutoFS, and Apache.

Fill out the form, then click the 'Next' button.

10. The public cluster network configuration screen allows you to set up the networking parameters for the ethernet network that connects the frontend to the outside network (e.g., the internet).

Welcome to Rocks

ROCKS

Help

Ethernet Configuration for Public Network

Public Network Interface:
This is the Ethernet network that physically connects your frontend to the outside world

Public Network Interface: eth1

IP address: 198.202.88.119

Netmask: 255.255.255.0

IP address:
Enter the IP address for eth1. This is the interface that connects the frontend to the outside network.

Netmask:
Enter the netmask for eth1:

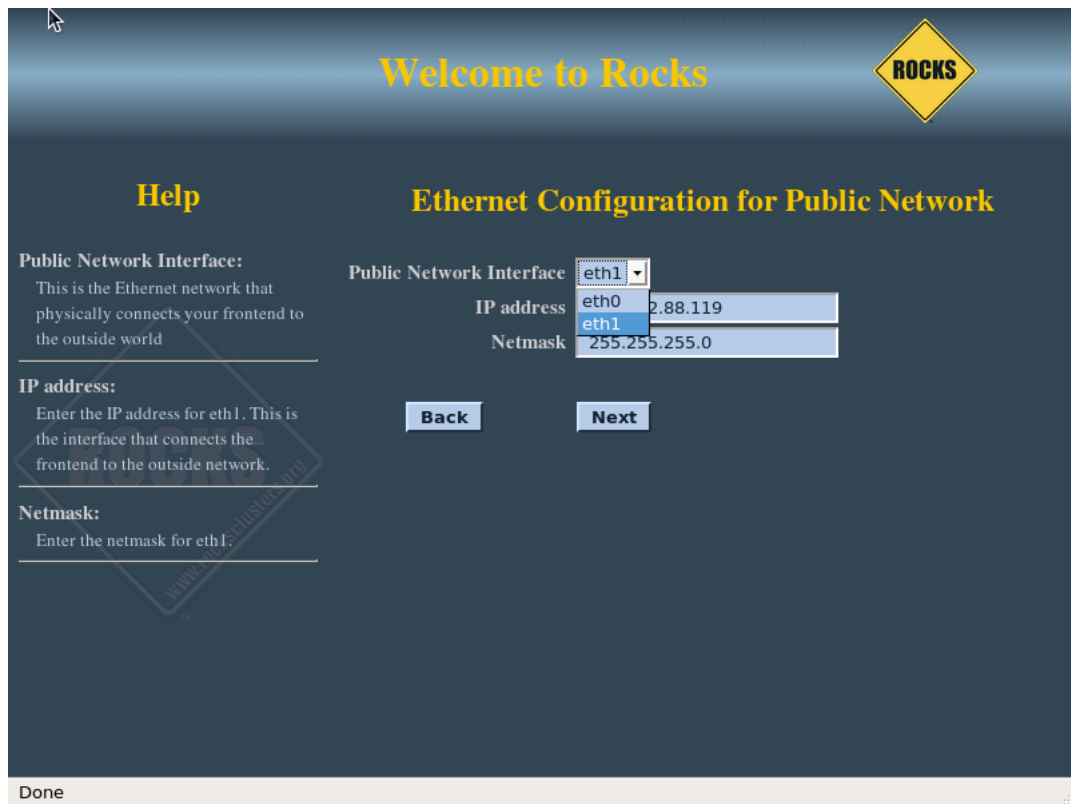
Back **Next**

Done

The above window is an example of how we configured the external network on one of our frontend machines.



The installer allows you select which physical interface is the public interface, if there is more than one interface. the network interface is a pull down menu as below



11. The private cluster network configuration screen allows you to set up the networking parameters for the ethernet network that connects the frontend to the compute nodes.

Welcome to Rocks

ROCKS

Help

Ethernet Configuration for Private Network

Private Network Interface:
This is the Ethernet network that physically connects your frontend to compute nodes

Private Network Interface: eth0

IP address: 10.1.1.1

Netmask: 255.255.0.0

IP address:
Enter the IP address for Private (cluster) Network. This is the interface that connects the frontend to the compute nodes.

Netmask:
Enter the netmask for private network.

Back **Next**

Done



It is recommended that you accept the defaults (by clicking the 'Next' button). But for those who have unique circumstances that requires different values for the internal ethernet connection, we have exposed the network configuration parameters.

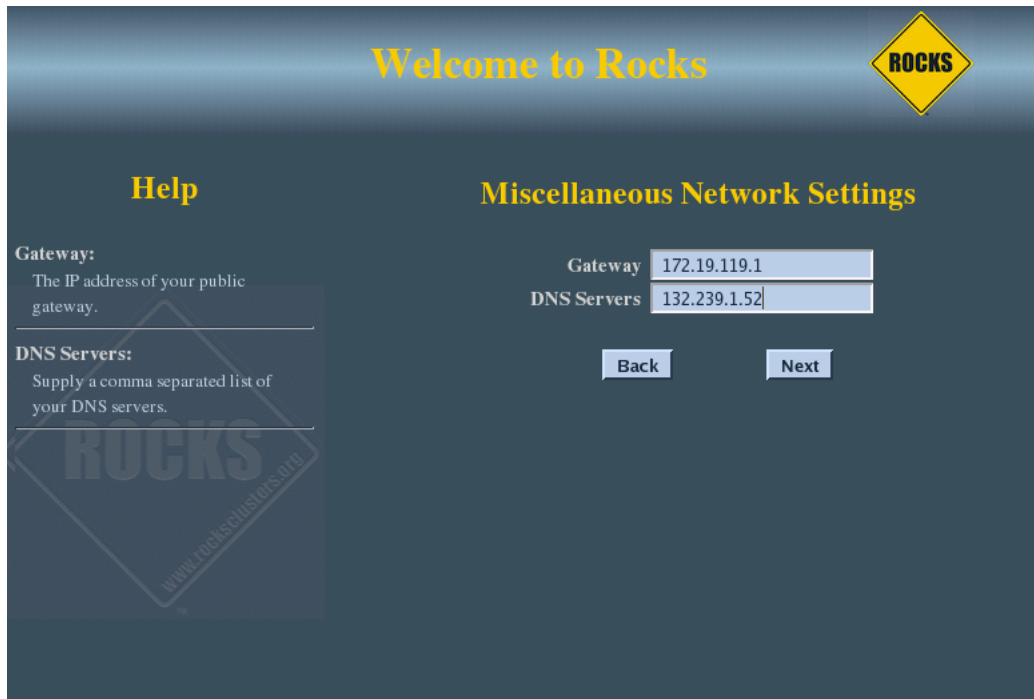


If you have only one physical interface, the installer will create a virtual ethernet interface (e.g. `eth0:0`).




The installer does not check if you selected the identical interface for your public and private interfaces. It is an error to do this

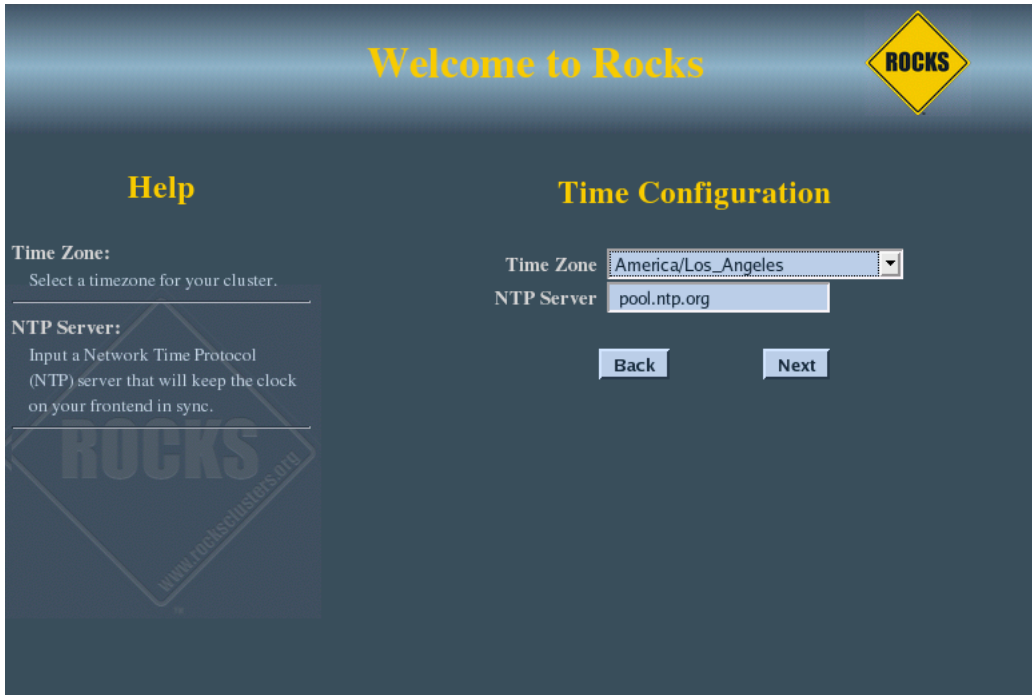
12. Configure the the *Gateway* and *DNS* entries:




13. Input the root password:



14. Configure the time:



15. The disk partitioning screen allows you to select *automatic* or *manual* partitioning.



To select automatic partitioning, click the `Auto Partitioning` radio button. This will repartition and reformat the first discovered hard drive that is connected to the frontend. All other drives connected to the frontend will be left untouched.

The first discovered drive will be partitioned like:

Table 3-1. Frontend -- Default Root Disk Partition

Partition Name	Size
----------------	------

Partition Name	Size
/	16 GB
/var	4 GB
swap	1 GB
/export (symbolically linked to /state/partition1)	remainder of root disk



When you use automatic partitioning, the installer will repartition and reformat the *first hard drive* that the installer discovers. All previous data on this drive will be erased. All other drives will be left untouched.

The drive discovery process uses the output of `cat /proc/partitions` to get the list of drives.

For example, if the node has an IDE drive (e.g., "hda") and a SCSI drive (e.g., "sda"), generally the IDE drive is the first drive discovered.

But, there are instances when a drive you don't expect is the first discovered drive (we've seen this with certain fibre channel connected drives). If you are unsure on how the drives will be discovered in a multi-disk frontend, then use manual partitioning.

16. If you selected manual partitioning, then you will now see Red Hat's manual partitioning screen:

The screenshot shows the manual partitioning interface for a drive labeled **Drive /dev/vda (35840 MB) (Model: Virtio Block Device)**. At the top, a table lists the drive's components:

/dev/vda1	16000 MB	/dev/vda2	4000 MB	/dev/vda5	14838 MB
-----------	----------	-----------	---------	-----------	----------

Below this is a table with columns: Device, Size (MB), Mount Point/ RAID/Volume, Type, and Format. The table lists the following partitions:

Device	Size (MB)	Mount Point/ RAID/Volume	Type	Format
vda1	16000	/	ext4	✓
vda2	4000	/var	ext4	✓
vda3	1000		swap	
vda4	14839		Extended	
vda5	14838	/export	ext4	✓

At the bottom of the window, there are buttons for **Create**, **Edit**, **Delete**, **Reset**, **Back**, and **Next**.

Above is an example of creating a '/', '/var', swap and '/export' partitions.




If you select manual partitioning, you must specify at least 16 GBs for the root partition and you must create a separate `/export` partition.



LVM is not supported by Rocks.

When you finish describing your partitions, click the 'Next' button.

17. The frontend will format its file systems, then it will ask for each of the roll CDs you added at the beginning of the frontend installation.



In the example screen above, insert the Kernel/Boot Roll into the CD tray and click 'OK'.


The contents of the CD will now be copied to the frontend's hard disk.

Repeat this step for each roll you supplied in steps 3-5.



After all the Rolls are copied, no more user interaction is required.

18. After the last roll CD is copied, the packages will be installed:




19. Finally, the boot loader will be installed and post configuration scripts will be run in the background. When they complete, the frontend will reboot.

3.4. Install Your Compute Nodes

1. Login to the frontend node as `root`.
2. Run the program which captures compute node DHCP requests and puts their information into the Rocks MySQL database:

```
# insert-ethers
```

This presents a screen that looks like:



If your frontend and compute nodes are connected via a managed ethernet switch, you'll want to select 'Ethernet Switches' from the list above. This is because the default behavior of many managed ethernet switches is to issue DHCP requests in order to receive an IP address that clients can use to configure and monitor the switch.


When `insert-ethers` captures the DHCP request for the managed switch, it will configure it as an ethernet switch and store that information in the MySQL database on the frontend.

As a side note, you may have to wait several minutes before the ethernet switch broadcasts its DHCP request. If after 10 minutes (or if `insert-ethers` has correctly detected and configured the ethernet switch), then you should quit `insert-ethers` by hitting the `F8` key.

Now, restart `insert-ethers` and continue reading below to configure your compute nodes.

Take the default selection, `Compute`, hit 'Ok'.

3. Then you'll see:



This indicates that `insert-ethers` is waiting for new compute nodes.

4. Power up the first compute node.



The BIOS boot order of your compute nodes should be: CD, PXE (Network Boot), Hard Disk.

If your compute nodes don't support PXE, then you'll need to boot your compute nodes with the Kernel Roll CD.

If you don't have a CD drive in your compute nodes and if the network adapters in your compute nodes don't support PXE, see Using a Floppy to PXE boot.

5. When the frontend machine receives the DHCP request from the compute node, you will see something similar to:

```
Insert Ethernet Addresses -- version 4.2
Opened kickstart access to 10.0.0.0/255.0.0.0 network

Inserted Appliances
Discovered New Appliance
Discovered a new appliance with MAC (00:13:72:ba:c8:df)

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

This indicates that `insert-ethers` received the DHCP request from the compute node, inserted it into the database and updated all configuration files (e.g., `/etc/hosts`, `/etc/dhcpd.conf` and DNS).

The above screen will be displayed for a few seconds and then you'll see the following:

```
Insert Ethernet Addresses -- version 4.2
Opened kickstart access to 10.0.0.0/255.0.0.0 network

Inserted Appliances
00:13:72:ba:c8:df      compute-0-0      ( )      #

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

In the above image, `insert-ethers` has discovered a compute node. The "()" next to `compute-0-0` indicates the node has not yet requested a kickstart file. You will see this type of output for each compute node that is successfully identified by `insert-ethers`.


```
rocks-85.sdsc.edu-restore-2017.11.16-0.x86_64.disk1.iso
```

Rocks 6

Burn your restore roll ISO image to a CD.

Rocks 6/7 Upgrade

Copy your restore roll ISO image to another machine for safekeeping.

Rocks 6

Reinstall the frontend by putting the Rocks Boot CD in the CD tray (generally, this is the Kernel/Boot Roll) and reboot the frontend.

Rocks 6/7 Upgrade

See full instructions for 6 to 7 upgrade. Restore roll is used after your Rocks 7 frontend has been installed

3.5.1. Upgrade From 6 to 7 using a Restore Roll Created on Rocks 6

Prior to upgrading to 7, you need to retain a copy of your restore roll iso image created using the process outlined in Section Upgrade Frontend.

Build your frontend as a fresh installation. See Install Frontend Rocks 6>. However, *do not reformat the partition that holds user home areas*. This is usually named `/export/home`

After installation of the frontend, you will need to copy your restore iso created in Section Upgrade Frontend to your newly installed system. For the rest of this section, we will use the example iso image name of `rocks-85.sdsc.edu-restore-2017.11.16-0.x86_64.disk1.iso`

After your normal frontend installation, there are some "fixups" that need to be applied to your restore roll. Once applied, the roll needs to be Added to your frontend and then "run". Finally, after the roll is run, some restored host keys need to be placed back into the rocks database.

The detailed instructions now follow

1. repack your existing restore roll

```
repack-roll rocks-85.sdsc.edu-restore-2017.11.17-0.x86_64.disk1.iso
```



this will create a copy of your restore roll with a *repacked.iso* ending. See the following screen output

```
# ls * iso
rocks-85.sdsc.edu-restore-2017.11.17-0.x86_64.disk1.iso
rocks-85.sdsc.edu-restore-2017.11.17-0.x86_64.disk1.repacked.iso
```

2. Add your repacked restore roll. Enable. Rebuild distro

```
# rocks add roll rocks-85.sdsc.edu-restore-2017.11.17-0.x86_64.disk1.repacked.iso
# rocks enable roll rocks-85.sdsc.edu-restore
# (cd /export/rocks/install; rocks create distro)
```

3. Run your roll

```
# rocks run roll rocks-85.sdsc.edu-restore | sh
```



You will see many warnings or errors. These are normal. The restore roll will not overwrite attributes that were entered/created during the frontend installation process. Some typical errors are shown in the screen below

```
Error: attribute "Kickstart_PublicHostname" exists
{attr} {value} [attr=string] [value=string]
Error: attribute "Kickstart_PublicAddress" exists
{attr} {value} [attr=string] [value=string]
...
Error: membership "Ethernet Switch" already exists
{appliance} [distribution=string] [graph=string] [membership=string] [node=string] [os=string]
...
Error: route exists
{host} {address} {gateway} [netmask=string]
```

4. Reload any keys that have been restored

```
# restore-keys
Removing host keys from Rocks DB and re-adding local...
Synching configuration and forcing make in /var/411
...
# su - root
# exit
```

5. Reinstall Compute nodes

```
# rocks set host boot compute action=install
# rocks run host compute "shutdown -r now"
```

3.5.2. Upgrade When Building a Rocks 6 Cluster

At the `boot :` prompt type:

```
build
```

At this point, the installation follows the same steps as a *normal* frontend installation (See the section: Install Frontend Rocks 6) -- with two exceptions:

1. On the first user-input screen (the screen that asks for 'local' and 'network' rolls), be sure to supply the Restore Roll that you just created.
2. You will be forced to manually partition your frontend's root disk.



You must reformat your `/` partition, your `/var` partition and your `/boot` partition (if it exists).

Also, be sure to assign the mountpoint of `/export` to the partition that contains the users' home areas. Do NOT erase or format this partition, or you will lose the user home directories. Generally, this is the largest partition on the first disk.

After your frontend completes its installation, the last step is to force a re-installation of all of your compute nodes. The following will force a PXE (network install) reboot of all your compute nodes.

```
# ssh-agent $SHELL
# ssh-add
# rocks run host compute '/boot/kickstart/cluster-kickstart-pxe'
```

3.5.3. Restore Roll Internals

By default, the Restore Roll contains two sets of files: system files and user files, and some user scripts. The system files are listed in the 'FILES' directive in the file:

```
/export/site-roll/rocks/src/roll/restore/src/system-files/version.mk.
FILES                = /etc/passwd /etc/shadow /etc/gshadow /etc/group \
                    /etc/exports /etc/auto.home /etc/motd
```

The user files are listed in the 'FILES' directive in the file:

```
/export/site-roll/rocks/src/roll/restore/version.mk.
FILES                += /etc/X11/xorg.conf
```

If you have other files you'd like saved and restored, then append them to the 'FILES' directive in the file `/export/site-roll/rocks/src/roll/restore/version.mk`, then rebuild the restore roll.

If you'd like to add your own post sections, you can add the name of the script to the 'SCRIPTS' directive of the `/export/site-roll/rocks/src/roll/restore/version.mk` file.

```
SCRIPTS += /share/apps/myscript.sh /share/apps/myscript2.py
```

This will add the shell script `/share/apps/myscript.sh`, and the python script `/share/apps/myscript2.py` in the post section of the `restore-user-files.xml` file.



If you'd like to run the script in "nochroot" mode, add

```
# nochroot
```

as the first comment in your script file after the interpreter line, if one is present.

For example

```
#!/bin/bash
#nochroot
echo "This is myscript.sh"
```

or

```
#nochroot
echo "This is myscript.sh"
```

will run the above code in the "nochroot" mode during installation. As opposed to

```
echo "This is myscript.sh"
#nochroot
```

or

```
#!/bin/bash
echo "This is myscript.sh"
```

will NOT run the script under "nochroot" mode.

All the files under `/export/rocks/install/site-profiles` are saved and restored. So, any user modifications that are added via the XML node method will be preserved.

The networking info for all node interfaces (e.g., the frontend, compute nodes, NAS appliances, etc.) are saved and restored. This is accomplished via the 'rocks dump' command.

3.6. Installing a Frontend over the Network


This section describes installing a Rocks frontend from a "Central" server over the wide area network, a process called WAN kickstart. The client frontend will retrieve Rocks Rolls and configuration over the Internet, and use these to install itself.

1. First, boot the node that will be your new frontend with the Kernel/Boot Roll CD (see steps 1 and 2 in the section "Install Frontend Rocks 6").
2. Then you'll see the screen as described in step 3 in the section "Install Frontend Rocks 6". Enter the FQDN of your central server in the *Hostname of Roll Server* text box (don't change this value if you want to use the default central server) then and click the `Download` button.

You'll see a screen that lists all the rolls available on the central server. Here's an example:




3. Now, select the rolls from the central server. To select a roll, click the checkbox next to roll. For example, this screen shows the *area51*, *base*, *bio* and *viz* rolls selected:



Click the *Submit* button to continue.

4. Now you'll see a screen similar to the screen below. This screen indicates that the *area51*, *base*, *bio* and *viz* rolls have been selected.



5. To select more rolls from another server, go to step 1 and enter a different FQDN.
6. If you'd like to include CD-based rolls with your Network-based rolls, click the *CD/DVD-based Roll* button and follow the instructions in the section "Install Frontend Rocks 6" starting at step 4.
7. When you are finished installing CD-based rolls, you will enter into the familiar Rocks installation windows. These may change depending on what rolls you have selected. Again the section "Install Frontend Rocks 6" has details for this process.
8. The installer will then retrieve the chosen rolls, rebuild the distribution with all rolls included, then install the packages. Finally, the installer will proceed with the post-section and other elements of a standard frontend install.

Your frontend should now be installed and ready to initialize compute nodes (see section Install Compute Nodes).

3.7. Enabling Public Web Access to Your Frontend

To permanently enable selected web access to the cluster from other machines on the public network, follow the steps below. Apache's access control directives will provide protection for the most sensitive parts of the cluster web site, however some effort will be necessary to make effective use of them.



HTTP (web access protocol) is a clear-text channel into your cluster. Although the Apache webserver is mature and well tested, security holes in the PHP engine have been found and exploited. Opening web access to the outside world by following the instructions below will make your cluster more prone to malicious attacks and breakins.

To open port 80 (the 'www' service) for the public network of frontend, execute:

```
# rocks remove firewall host=localhost rulename=A40-WWW-PUBLIC-LAN
# rocks add firewall host=localhost network=public protocol=tcp service=www chain=INPUT \
  action=ACCEPT flags="-m state --state NEW --source 0.0.0.0/0.0.0.0" \
  rulename=A40-WWW-PUBLIC-NEW
```

Then we can see the what the resulting firewall rules will look like:

```
# rocks report host firewall localhost
```

```

<file name="/etc/sysconfig/iptables" perms="500">
*nat
# MASQUERADE (host) :
-A POSTROUTING -o eth1 -j MASQUERADE
COMMIT

*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
# A10-REJECT-411-TCP (host) :
-A INPUT -p tcp --dport 372 --sport 1024:65535 -j REJECT
# A10-REJECT-411-UDP (host) :
-A INPUT -p udp --dport 372 --sport 1024:65535 -j REJECT
# A15-ALL-LOCAL (global) :
-A INPUT -i lo -j ACCEPT
# A20-ALL-PRIVATE (global) :
-A INPUT -i eth0 -j ACCEPT
# A20-SSH-PUBLIC (global) :
-A INPUT -i eth1 -p tcp --dport ssh -m state --state NEW -j ACCEPT
# A30-RELATED-PUBLIC (global) :
-A INPUT -i eth1 -m state --state RELATED,ESTABLISHED -j ACCEPT
# A40-HTTPS-PUBLIC-LAN (host) :
-A INPUT -i eth1 -p tcp --dport https -m state --state NEW --source &Kickstart_PublicNetwork;/&Ki
# A40-WWW-PUBLIC-NEW (host) :
-A INPUT -i eth1 -p tcp --dport www -m state --state NEW --source 0.0.0.0/0.0.0.0 -j ACCEPT
# A50-FORWARD-RELATED (host) :
-A FORWARD -i eth1 -o eth0 -m state --state RELATED,ESTABLISHED -j ACCEPT
# A60-FORWARD (host) :
-A FORWARD -i eth0 -j ACCEPT
# R10-GANGLIA-UDP (host) : block ganglia traffic from non-private interfaces
-A INPUT -p udp --dport 8649 -j REJECT
# R20-MYSQL-TCP (host) : block mysql traffic from non-private interfaces
-A INPUT -p tcp --dport 3306 -j REJECT
# R30-FOUNDATION-MYSQL (host) : block foundation mysql traffic from non-private interfaces
-A INPUT -p tcp --dport 40000 -j REJECT
# R900-PRIVILEGED-TCP (global) :
-A INPUT -i eth1 -p tcp --dport 0:1023 -j REJECT
# R900-PRIVILEGED-UDP (global) :
-A INPUT -i eth1 -p udp --dport 0:1023 -j REJECT
COMMIT
</file>

```

In the above example, eth0 is associated with the private network and eth1 is associated with the public network.

Notice the line: "-A INPUT -i eth1 -p tcp --dport www -m state --state NEW --source 0.0.0.0/0.0.0.0 -j ACCEPT". This is the line in the firewall configuration that will allow web traffic from any source to flow in and out of the frontend. This line was added to your firewall configuration with the "rocks add firewall host=localhost" command that you executed.

Also, notice that the original line was: "-A INPUT -i eth1 -p tcp --dport www -m state --state NEW --source &Kickstart_PublicNetwork;/&Kickstart_PublicNetmask; -j ACCEPT". This default Rocks firewall rule allows web traffic from *your local public subnet* to flow in and out of the frontend.

Now apply the configuration to the host:

```
# rocks sync host firewall localhost
```

The host will now accept web traffic on its public interface.

Test your changes by pointing a web browser to `http://my.cluster.org/`, where "my.cluster.org" is the DNS name of your frontend machine.



If you cannot connect to this address, the problem is most likely in your network connectivity between your web browser and the cluster. Check that you can ping the frontend machine from the machine running the web browser, that you can ssh into it, etc.

Notes

1. <http://www.oreilly.com>
2. <https://wiki.centos.org/HowTos/InstallFromUSBKey>

Chapter 4. Defining and Modifying Networks and Network Interfaces

4.1. Networks, Subnets, VLANs and Interfaces

Rocks uses a SQL database to hold information about nodes including network device information. In version 5.1 support was added for VLAN tagged interfaces to enable construction of virtual clusters and other more complicated network scenarios. There are a large number of commands that allow manipulation of subnet definitions, physical interfaces, and logical VLAN interfaces.



The basic model of operation is for an administrator to use a series of commands to add and set/change networking definitions in the database and then either re-install a node or reconfigure/restart the network configuration by calling `rocks sync config <host>`.

4.2. Named Networks/Subnets

Rocks clusters are required to have two subnets defined: "public" and "private", but a cluster owner can define more subnets. The command `rocks list network` lists the defined networks

```
[root@rocks ~]# rocks list network
NETWORK      SUBNET      NETMASK
private:     172.16.254.0 255.255.255.0
public:      132.239.8.0  255.255.255.0
optiputer:   67.58.32.0   255.255.224.0
```

In the screen above, the additional network called "optiputer" is defined with netmask 255.255.224.0(/19). To add a network called "fast" as 192.168.1.0 and netmask 255.255.255.0(/24) do the following

```
[root@rocks ~]# rocks add network fast subnet=192.168.1.0 netmask=255.255.255.0
[root@rocks ~]# rocks list network
NETWORK      SUBNET      NETMASK
private:     172.16.254.0 255.255.255.0
public:      132.239.8.0  255.255.255.0
optiputer:   67.58.32.0   255.255.224.0
fast:        192.168.1.0  255.255.255.0
```

The subnet and netmask of an existing network can be changed using `rocks set network subnet` and `rocks set network netmask` commands.

4.3. Host Interfaces

There are three types of interfaces that a cluster owner may need to be concerned about: physical, logical, and VLAN (virtual LAN) bridges. Linux (and other OSes like Solaris) support logical interfaces that share a particular physical network port. The following shows physical network devices and associations of those devices to a named network (or subnet, used interchangeably in this discussion). In the figures below, the `<nn>`

notation is a standard method of how to specify the number of bits in the netmask. Examples include: /24=255.255.255.0 (Class C subnet), /16=255.255.0.0 (Class B subnet), /8=255.0.0.0 (Class A subnet) and /25=255.255.255.128




FIGURE: Hosts can have any number of physical networking devices. Every Rocks node must have a private network defined (e.g., eth0). Frontends also must have a separate public network (e.g., eth1). Other devices could be myri0 (for Myrinet) or ib0 (for Infiniband).

Adding a new network interface to a host can be done from the command line. For example, to add an interface named "myri0" with IP address 192.168.1.10 on the logical subnet "fast":

```
[root@rocks ~]# rocks add host interface compute-0-0-1 iface=myri0 subnet=fast ip=192.168.1.10
[root@rocks ~]# rocks list host interface compute-0-0-1
```

SUBNET	IFACE	MAC	IP	NETMASK	MODULE NAME	VLAN
private	eth0	00:16:3e:00:00:11	172.16.254.192	255.255.255.0	xennet compute-0-0-1	-----
fast	myri0	-----	192.168.1.10	255.255.255.0	-----	-----

You can also set other fields for a host interface (if the field is one of [mac, ip, module, name, vlan]) with the command `rocks set host interface <field> <host> iface=<iface> value`. To set the name associated with the myri0 interface to compute-myri-0-0-1 on the node compute-0-0-1, execute:

```
[root@rocks ~]# rocks set host interface name compute-0-0-1 iface=myri0 compute-myri-0-0-1
[root@rocks ~]# rocks list host interface compute-0-0-1
```

SUBNET	IFACE	MAC	IP	NETMASK	MODULE NAME	VLAN
private	eth0	00:16:3e:00:00:11	172.16.254.192	255.255.255.0	xennet compute-0-0-1	-----
fast	myri0	-----	192.168.1.10	255.255.255.0	compute-myri-0-0-1	-----

4.4. Virtual LANs (VLANs) and Logical VLAN Bridges

Linux supports VLAN tagging on virtual interfaces (i.e., IEEE 802.1Q). For example, if a host has physical interface eth0 (untagged), then the kernel can send and receive a tagged packets if a properly defined interface named eth0.<vlan> has been created and properly configured. Tagging allows the same physical network to be partitioned into many different networks. A key feature of VLAN tagging is that a broadcast packet (e.g. a DHCPDISCOVER packet) only broadcasts on the tagged VLAN in which it was initially sent.

Rocks supports two types of VLAN interfaces - the first is an explicit device name like eth0.10 that is defined on a particular physical interface. The second is a logical device name of the form "vlan*". In Rocks, the physical VLAN device can also have an IP address associated with it, however a logical VLAN device cannot. We use logical VLANs to construct bridges suitable for virtual clusters.



1. Explicit VLAN Devices of the form <interface>.<vlan> can have IP addresses assigned
2. Rocks-Specific: Logical VLAN Devices of the form "vlan*" CANNOT have IP address assigned

4.4.1. Physical VLAN Devices

Physical VLAN devices are interfaces associated with specific physical interfaces. While eth0 is used as an example, any physical IP interface can have a VLAN associated with it.




FIGURE: Physical VLAN device called eth0.2. This device may be assigned an IP and a network name (e.g. "net") that is unrelated to the network name of the physical device (eth0). All packets sent on this interface will be tagged with VLAN=2. Multiple Physical VLAN devices can be defined.

Use the following example to add a physical VLAN device, assign a tag, and an IP address:

```
[root@rocks ~]# rocks add host interface compute-0-0-1 iface=eth0.2 subnet=net2 ip=10.2.1.10
[root@rocks ~]# rocks set host interface vlan compute-0-0-1 iface=eth0.2 vlan=2
[root@rocks ~]# rocks list host interface compute-0-0-1
```

SUBNET	IFACE	MAC	IP	NETMASK	MODULE	NAME	VLAN
private	eth0	00:16:3e:00:00:11	172.16.254.192	255.255.255.0	xennet	compute-0-0-1	-----
net2	eth0.2	-----	10.2.1.10	255.255.255.0	-----	-----	2

4.4.2. Logical VLAN Devices

The second kind of VLAN interface that Rocks supports is what we call a *logical VLAN device*. The Virtual VLAN gives the ability of having a raw interface with no IP address assigned that is generally used as a bridge for virtual machines. Virtual VLAN devices have their subnet=<subnet of physical>




FIGURE: Virtual VLAN devices called vlan2 and vlan3. These types of devices may NOT have an IP address (This is a Rocks-specific construction).

```
[root@rocks ~]# rocks add host interface compute-0-0-1 vlan2
[root@rocks ~]# rocks add host interface compute-0-0-1 vlan3
[root@rocks ~]# rocks set host interface vlan compute-0-0-1 vlan2 2
[root@rocks ~]# rocks set host interface vlan compute-0-0-1 vlan3 3
[root@rocks ~]# rocks list host interface compute-0-0-1
```

SUBNET	IFACE	MAC	IP	NETMASK	MODULE	NAME	VLAN
private	eth0	00:16:3e:00:00:11	172.16.254.192	255.255.255.0	xennet	compute-0-0-1	-----
-----	vlan2	-----	-----	-----	-----	-----	2
-----	vlan3	-----	-----	-----	-----	-----	3

At this stage, the vlan interfaces are not currently associated with any physical network device. Linux will not configure these devices on the node without the association. We overload the meaning of subnet in this case to mean: "associate the logical vlan device with the physical device that is in subnet 'x'". As an example, we can associate both vlan2 and vlan3 to be tagged packet interfaces on the the subnet named private.

```
[root@tranquil ~]# rocks set host interface subnet compute-0-0-1 vlan2 subnet=private
[root@tranquil ~]# rocks set host interface subnet compute-0-0-1 vlan3 subnet=private
[root@tranquil ~]# rocks list host interface compute-0-0-1
```

SUBNET	IFACE	MAC	IP	NETMASK	MODULE	NAME	VLAN
private	eth0	00:16:3e:00:00:11	172.16.254.192	255.255.255.0	xennet	compute-0-0-1	-----
private	vlan2	-----	-----	-----	-----	-----	2
private	vlan3	-----	-----	-----	-----	-----	3

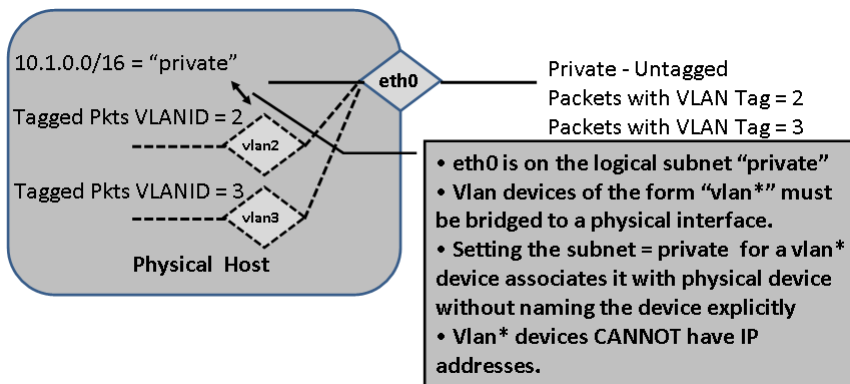


FIGURE: Virtual VLAN devices called vlan2 and vlan3 are associated with the physical device that is designated as subnet private. Notice, that no netmask is associated with the vlan2 and vlan3 devices. These are raw, tagged packet interfaces and are mostly used for bridges when hosting VMs.

4.5. Configuring IPMI interfaces

The Intelligent Platform Management Interface (IPMI) is a standardized computer system interface used by system administrators to manage a computer system and monitor its operation. IPMI can be used to remotely power cycle a server, remotely connect to its KVM or serial interface and perform other administrative related functionalities. The standard is supported by most hardware vendor (Dell, HP, NEC, Intel, SuperMicro, etc.) and it is found on several server platforms under different name (ILOM, BMC, DRAC, etc.).

Rocks by default include openIPMI¹ an Open Source command line implementation of the standard on all the nodes and it is possible to create the initial configuration of the devices leveraging Rocks database and Rocks command line.

4.5.1. Configuring a IPMI Device

Rocks uses its data base to keep configuration information regarding the various IPMI devices. All interfaces defined in the rocks database named "ipmi" are treated as IPMI devices and not as standard ethernet devices.

It is good practice to define a network to allocate all the IPMI devices (although not necessary). IPMI network are sometimes separated from other network for security purposes, other times they can share the existing network infrastructure to save money. The network configuration and design is behind the scope of this tutorial. We simply assume All the nodes are connected to the admin network (10.10.0.0/255.255.0.0) and the frontend has some routing access to it. In the following example can be used to define the admin network in the rocks DB.

```
# rocks add network admin subnet=10.10.0.0 netmask=255.255.0.0 mtu=1500
```

To configure a IPMI device on a node called compute-0-0, you should run:

```
# rocks add host interface compute-0-0 ipmi ip=10.10.255.254 name=ipmi-compute-0-0 subnet=admin
# rocks set host interface channel compute-0-0 ipmi 1
```

The channel is used to specify which communication interface should be configured. Each IPMI device supports multiple communication devices (Ethernet/IP, serial ports, etc.) so it is necessary to identify which one is associated with the Ethernet interface to specify it with the rocks set host interface channel command. Different hardware can have different channels for the same type of device. To this purpose it is possible to use ipmitool command repeatedly with different channel numbers, untill the Medium info display 802 LAN, as in the following example:

```
# ipmitool channel info 1
Channel 0x1 info:
Channel Medium Type      : 802.3 LAN
Channel Protocol Type    : IPMB-1.0
...
```

The admin password is set by default to "admin", it is possible to change it using the rocks attribute "ipmi_password". In most of the IPMI devices the admin user ID is equal to 1. If a device set up in a different way it is possible to set the attribute "ipmi_admin_user_number" and change this default behavior.

The command rocks sync network compute-0-0-0 will create a script on compute-0-0-0 in /etc/sysconfig/ipmi which should be executed manually on the node to finalize the configuration.

```
# rocks sync host network compute-0-0
# ssh compute-0-0
# bash /etc/sysconfig/ipmi
```

After this your IPMI interface should be properly configured. To validate your configuration you can use the ipmitool command from your frontend (or from the node which has access to the admin network) as shown in the following example:

```
ipmitool -H 10.10.255.254 -U ADMIN -P admin power status
```

To power off compute-0-0, execute:

```
# ipmitool -H 10.10.255.254 -U ADMIN -P admin power off
```

To power compute-0-0 back on, execute:

```
# ipmitool -H 10.10.255.254 -U ADMIN -P admin power on
```

There are several more management tasks you can perform with ipmitool. See the man page for ipmitool for more details.

4.6. Networking Configuration Examples

In this section, we describe some common networking configurations and how to use Rocks commands to set up various networking scenarios.

4.6.1. Adding a public IP address to the second ethernet adapter on a compute node

Often, owners want the second ethernet adapter to be on the public network and for the default routing to be in the public network. Assuming that the public network is 1.2.3.0/255.255.255.0 and the default gateway for that network is 1.2.3.1, the following set of commands define the second interface of a compute to have address 1.2.3.25 with name mypublic.<public domain of frontend> update all required configuration files on the frontend, update all required configuration files on the node compute-0-0 and restart the network on compute-0-0.

```
# rocks set host interface ip compute-0-0 iface=eth1 ip=1.2.3.25
# rocks set host interface name compute-0-0 iface=eth1 name=mypublic
# rocks set host interface subnet compute-0-0 eth1 public
# rocks add host route compute-0-0 0.0.0.0 1.2.3.1 netmask=0.0.0.0
# rocks sync config
# rocks sync host network compute-0-0
```

4.6.2. Adding an IP network for local message passing.

Often, users will want to use the second ethernet device for messaging passing. In this example, we illustrate creating a named subnet and then scripting the IP assignment for a rack of 32 nodes with IP range of 192.168.1.10 ... 192.168.1.41.

```
rocks add network fast subnet=192.168.1.0 netmask=255.255.255.0
IP=10
NNODES=32
NODE=0
while [ $NODE -lt $NNODES ]; do \
  rocks set host interface ip compute-0-$NODE iface=eth1 ip=192.168.1.$IP; \
  rocks set host interface subnet compute-0-$NODE iface=eth1 subnet=fast; \
  rocks set host interface name compute-0-$NODE iface=eth1 name=compute-fast-0-$NODE; \
  let IP++; \
  let NODE++; \
done
rocks sync config
rocks sync host network compute
```

The above will add the named subnet called "fast", assign IP addresses sequentially, name the eth1 interface on each node, rewrite the DNS configuration (sync config) and finally rewrite and restart the network configuration on each compute appliance. This additional network configuration is persistent across re-installation of nodes.

4.6.3. Setting the Default Gateway on a Node

When adding a second (public) interface on nodes, it is sometimes desirable to change the default route or gateway. Suppose that your gateway=1.2.3.4, the the following will set gateway on a particular host

```
rocks add host route compute-0-0 0.0.0.0 1.2.3.4 netmask=0.0.0.0
rocks sync host network compute-0-0
```



the 1.2.3.4 address must be directly reachable (Same LAN) from one of the host's ethernet interfaces.

Notes

1. <http://openipmi.sourceforge.net/>

Chapter 5. Customizing your Rocks Installation

5.1. Adding Packages to Compute Nodes

Put the package you want to add in:

```
/export/rocks/install/contrib/7.0/arch/RPMS
```

Where *arch* is your architecture ("i386" or "x86_64").

Create a new XML configuration file that will *extend* the current `compute.xml` configuration file:

```
# cd /export/rocks/install/site-profiles/7.0/nodes
# cp skeleton.xml extend-compute.xml
```



If you use `extend-compute.xml` your packages will be installed only on your computed nodes. If you want your packages to be installed on all other appliances (e.g. login nodes, nas nodes, etc.) you should use `extend-base.xml` instead of `extend-compute.xml`.

Inside `extend-compute.xml`, add the package name by changing the section from:

```
<package> <!-- insert your package name here --> </package>
```

to:

```
<package> your package </package>
```



It is important that you enter the *base name* of the package in `extend-compute.xml` and not the full name.

For example, if the package you are adding is named `XFree86-100dpi-fonts-4.2.0-6.47.i386.rpm`, input `XFree86-100dpi-fonts` as the package name in `extend-compute.xml`.

```
<package>XFree86-100dpi-fonts</package>
```

If you have multiple packages you'd like to add, you'll need a separate `<package>` tag for each. For example, to add both the 100 and 75 dpi fonts, the following lines should be in `extend-compute.xml`:

```
<package>XFree86-100dpi-fonts</package>
<package>XFree86-75dpi-fonts</package>
```

Also, make sure that you remove any package lines which do not have a package in them. For example, the file should NOT contain any lines such as:

```
<package> <!-- insert your package name here --> </package>
```

Now build a new Rocks distribution. This will bind the new package into a RedHat compatible distribution in the directory `/export/rocks/install/rocks-dist/...`

```
# cd /export/rocks/install
# rocks create distro
```

Now, reinstall your compute nodes.

5.1.1. Adding Specific Architecture Packages to Compute Nodes

Often on `x86_64`-based clusters, one wants to add the `x86_64` and `i386` version of a package to compute nodes. To do this, in your `extend-compute.xml` file, supply the section:

```
<package>pkg.x86_64</package>
<package>pkg.i386</package>
```

Where `pkg` is the basename of the package.

Now build a new Rocks distribution.

```
# cd /export/rocks/install
# rocks create distro
```

Now, reinstall your compute nodes.

5.2. Customizing Configuration of Compute Nodes

Create a new XML configuration file that will *extend* the current `compute.xml` configuration file:

```
# cd /export/rocks/install/site-profiles/7.0/nodes/
# cp skeleton.xml extend-compute.xml
```



If you want to include your customizations in all cluster nodes (and not just the compute nodes) you should use `extend-base.xml` as a file name instead of `extend-compute.xml`.

Inside `extend-compute.xml`, add your configuration scripts that will be run in the *post configuration* step of the Red Hat installer.

Put your bash scripts in between the tags `<post>` and `</post>`:

```
<post>
<!-- insert your scripts here -->
</post>
```

To apply your customized configuration scripts to compute nodes, rebuild the distribution:

```
# cd /export/rocks/install
# rocks create distro
```

Then, reinstall your compute nodes.

5.3. Adding Applications to Compute Nodes

If you have code you'd like to share among the compute nodes, but your code isn't in an RPM (or in a roll), then this procedure describes how you can share it with NFS.

On the frontend, go to the directory `/share/apps`.

```
# cd /share/apps
```

Then add the files you'd like to share within this directory.

All files will also be available on the compute nodes under: `/share/apps`. For example:

```
# cd /share/apps
# touch myapp
# ssh compute-0-0
# cd /share/apps
# ls
myapp
```

5.4. Configuring Additional Ethernet Interfaces

For compute nodes, Rocks uses the first ethernet interface (`eth0`) for management (e.g., reinstallation), monitoring (e.g., Ganglia) and message passing (e.g., OpenMPI over ethernet). Often, compute nodes have more than one ethernet interface. This procedure describes how to configure them.

Additional ethernet interfaces are configured from the frontend via the Rocks command line. It modifies entries in the networks table on the frontend to add information about an extra interface on a node.

Once you have the information in the networks table, every time you reinstall, the additional NIC will be configured.

Suppose you have a compute node with one configured network (`eth0`) and one unconfigured network (`eth1`):

```
# rocks list host interface compute-0-0
SUBNET  IFACE  MAC                IP                NETMASK          MODULE  NAME                VLAN
private eth0    00:1e:4f:b0:74:ef  10.1.255.254     255.255.0.0      tg3            compute-0-0  -----
----- eth1    00:10:18:31:74:43  -----          -----          tg3            -----      -----
```

We'll configure `eth1` with the following network info and associate `eth1` with the public subnet:

- Name = `fast-0-0`
- IP address = `192.168.1.1`

```
# rocks set host interface ip compute-0-0 eth1 192.168.1.1
# rocks set host interface name compute-0-0 eth1 fast-0-0
```

Now we'll create a new 'network' and associate it with the new interface:

```
# rocks add network fast 192.168.1.0 255.255.255.0
```

And then we'll check our work:

```
# rocks list network
NETWORK  SUBNET          NETMASK          MTU
private: 10.1.0.0      255.255.0.0     1500
public:  137.110.119.0 255.255.255.0  1500
fast:    192.168.1.0   255.255.255.0  1500
```

Now associate the new network to eth1.

```
# rocks set host interface subnet compute-0-0 eth1 fast
```

The interface eth1 is now configured:

```
# rocks list host interface compute-0-0
SUBNET  IFACE  MAC          IP          NETMASK      MODULE  NAME          VLAN
private eth0   00:1e:4f:b0:74:ef 10.1.255.254 255.255.0.0 tg3         compute-0-0  -----
fast    eth1   00:10:18:31:74:43 192.168.1.1  255.255.255.0 tg3         fast-0-0     -----
```

After specifying new network settings to a compute-0-0, execute the following command to apply the settings:

```
# rocks sync config
# rocks sync host network compute-0-0
```



If you configuring the interface to another public network, you can set the gateway for the interface with the `rocks add host route` command.

For example, to set the route for the 192.168.1.0 network to 192.168.1.254 for compute-0-0, you'd execute:

```
# rocks add host route compute-0-0 192.168.1.0 192.168.1.254 netmask=255.255.255.0
```

5.5. Compute Node Disk Partitioning

5.5.1. Default Disk Partitioning

The default root partition is 16 GB, the default swap partition is 1 GB, and the default /var partition is 4 GB. The remainder of the root disk is setup as the partition `/state/partition1`.

Only the root disk (the first discovered disk) is partitioned by default. To partition all disks connected to a compute node, see the section [Forcing the Default Partitioning Scheme for All Disks on a Compute Node](#).

Table 5-1. Compute Node -- Default Root Disk Partition

Partition Name	Size
/	16 GB

Partition Name	Size
swap	1 GB
/var	4 GB
/state/partition1	<i>remainder of root disk</i>



After the initial installation, all data in the file systems labeled `/state/partitionX` will be preserved over reinstallations.

5.5.2. Customizing Compute Node Disk Partitions

In Rocks, to specify custom partitioning for a node, one must write code in a `<pre>` section which creates a file named `/tmp/user_partition_info`. with Red Hat kickstart partitioning directives in it. This file will be used to create the initial partitioning on a node when no other info is contained in the Frontend data base and the pre-existing partition are not marked by rocks (e.g. `./rocks-release`). While when reinstalling a node which has partition information already stored in the data base and the local disk are already marked Rocks installer will simply reformat the `"/"` partition and the `"/var"` and leave untouched all the other partitions. This allows users to fully program their cluster nodes' partitions. In the examples below, we'll explore what this means.

In Rocks 7, the `<pre>` section of any custom partitioning must be executed prior to `<pre>` section of the node `partition.xml`. In the default configuration graph the node `custom-partition` satisfies this constraint. The simplest place to put your custom partitioning code is to replace this node file.

5.5.2.1. Single Disk Example

Create a new XML node file that will *replace* the current `custom-partition.xml` XML node file:

```
# cd /export/rocks/install/site-profiles/7.0/nodes/
# cp /export/rocks/install/rocks-dist/x86_64/build/nodes/custom-partition.xml replace-custom-part
```

Inside `replace-custom-partition.xml`, add the following section right after the `<main>` `</main>` section:

```
<main>
  <!-- kickstart 'main' commands go here -->
</main>

<pre>
echo "clearpart --all --initlabel --drives=hda
part / --size 8000 --ondisk hda
part swap --size 1000 --ondisk hda
part /mydata --size 1 --grow --ondisk hda" &gt; /tmp/user_partition_info
</pre>
```



Important! you need to use the XML code for the `">"` in your node file. That code is `"&>"`

The above example uses a bash script to populate `/tmp/user_partition_info`. This will set up an 8 GB root partition, a 1 GB swap partition, and the remainder of the drive will be set up as `/mydata`. Additional drives on your compute nodes can be setup in a similar manner by changing the `--ondisk` parameter.

In the above example, the syntax of the data in `/tmp/user_partition_info` follows directly from Red Hat's kickstart. For more information on the `part` keyword, see Red Hat Enterprise Linux 5 Installation Guide : Kickstart Options¹.



User-specified partition mountpoint names (e.g., `/mydata`) cannot be longer than 15 characters.

Then apply this configuration to the distribution by executing:

```
# cd /export/rocks/install
# rocks create distro
```

To reformat compute node `compute-0-0` to your specification above, you'll need to first remove the partition info for `compute-0-0` from the database:

```
# rocks remove host partition compute-0-0
```

Then you'll need to remove the file `.rocks-release` from the first partition of *each disk* on the compute node. Here's an example script:

```
for file in $(mount | awk '{print $3}')
do
  if [ -f $file/.rocks-release ]
  then
    rm -f $file/.rocks-release
  fi
done
```

Save the above script as `/share/apps/nukeit.sh` and then execute:

```
# ssh compute-0-0 'sh /share/apps/nukeit.sh'
```

Then, reinstall the node:

```
# rocks set host boot action=install compute-0-0
# ssh compute-0-0 'shutdown -r now'
```

5.5.2.2. Software Raid Example

If you would like to use software RAID on your compute nodes, inside `replace-custom-partition.xml` add a section that looks like:

```
<pre>
echo "clearpart --all --initlabel --drives=hda,hdb
part / --size 8000 --ondisk hda
part swap --size 1000 --ondisk hda

part raid.00 --size=10000 --ondisk hda
part raid.01 --size=10000 --ondisk hdb

raid /mydata --level=1 --device=md0 raid.00 raid.01" > /tmp/user_partition_info
```

```
</pre>
```

Then apply this configuration to the distribution by executing:

```
# cd /export/rocks/install
# rocks create distro
```

To reformat compute node `compute-0-0` to your specification above, you'll need to first remove the partition info for `compute-0-0` from the database:

```
# rocks remove host partition compute-0-0
```

Then you'll need to remove the file `.rocks-release` from the first partition of *each disk* on the compute node. Here's an example script:

```
for file in $(mount | awk '{print $3}')
do
  if [ -f $file/.rocks-release ]
  then
    rm -f $file/.rocks-release
  fi
done
```

Save the above script as `/share/apps/nukeit.sh` and then execute:

```
# ssh compute-0-0 'sh /share/apps/nukeit.sh'
```

Then, reinstall the node:

```
# rocks set host boot action=install compute-0-0
# ssh compute-0-0 'shutdown -r now'
```

5.5.2.3. Programmable Partitioning

Some issues with the above two examples are that 1) you must know the name of the disk device (e.g., `hda`) and, 2) the partitioning will be applied to all nodes. We can avoid these issues by writing a python program that emits node-specific partitioning directives.

In the next example, we'll use some Rocks partitioning library code to dynamically determine the name of the boot disk.

```
<pre arg="--interpreter /opt/rocks/bin/python">
```

```
from rocks import rocks_partition

membership = '&membership;'
nodename = '&hostname;'

def doDisk(file, disk):
    file.write('clearpart --all --initlabel --drives=%s\n' % disk)
    file.write('part / --size=6000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part /var --size=2000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part swap --size=2000 --ondisk=%s\n' % disk)
    file.write('part /mydata --size=1 --grow --fstype=ext3 --ondisk=%s\n'
               % disk)

#
```

```
# main
#
p = rocks_partition.RocksPartition()
disks = p.getDisks()

if len(disks) == 1:
    file = open('/tmp/user_partition_info', 'w')
    doDisk(file, disks[0])
    file.close()
</pre>
```

The function `getDisks()` returns a list of discovered disks. In the code sample above, if only one disk is discovered on the node, then the function `doDisk` is called which outputs partitioning directives for a single disk. This code segment will work for nodes with IDE or SCSI controllers. For example, a node with a IDE controller will name its disks `hdX` and a node with SCSI controllers will name its disks `sdX`. But, the code segment above doesn't care how the node names its drives, it only cares if one drive is discovered.

The next example shows how a node can automatically configure a node for software raid when it discovers 2 disks. But, if the node only discovers 1 disk, it will output partitioning info appropriate for a single-disk system.

```
<pre arg="--interpreter /opt/rocks/bin/python">

from rocks import rocks_partition

membership = '&membership;'
nodename = '&hostname;'

def doRaid(file, disks):
    file.write('clearpart --all --initlabel --drives=%s\n'
              % ','.join(disks))

    raidparts = []

    for disk in disks:
        if disk == disks[0]:
            part = 'part / --size=6000 --fstype=ext3 ' + \
                  '--ondisk=%s\n' % disk
            file.write(part)

            part = 'part /var --size=2000 --fstype=ext3 ' + \
                  '--ondisk=%s\n' % disk
            file.write(part)

            part = 'part raid.%s --size=5000 --ondisk=%s\n' % (disk, disk)
            file.write(part)

            raidparts.append('raid.%s' % disk)

    raid = 'raid /bigdisk --fstype=ext3 --device=md0 --level=1 %s\n' \
          % ','.join(raidparts)
    file.write(raid)

def doDisk(file, disk):
    file.write('clearpart --all --initlabel --drives=%s\n' % disk)
    file.write('part / --size=6000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part /var --size=2000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part swap --size=2000 --ondisk=%s\n' % disk)
```

```

        file.write('part /mydata --size=1 --grow --fstype=ext3 --ondisk=%s\n'
                  % disk)

#
# main
#
p = rocks_partition.RocksPartition()
disks = p.getDisks()

file = open('/tmp/user_partition_info', 'w')

if len(disks) == 2:
    doRaid(file, disks)
elif len(disks) == 1:
    doDisk(file, disks[0])

file.close()

</pre>

```

If the node has 2 disks (if `len(disks) == 2:`), then call `doRaid()` to configure a software raid 1 over the 2 disks. If the node has 1 disk then call `doDisk()` and output partitioning directives for a single disk.

In the next example, we show how to output user-specified partitioning info for only one specific node (`compute-0-0`). All other nodes that execute this pre section will get the default Rocks partitioning.

```

<pre arg="--interpreter /opt/rocks/bin/python">

from rocks import rocks_partition

membership = '&membership;'
nodename = '&hostname;'

def doRaid(file, disks):
    file.write('clearpart --all --initlabel --drives=%s\n'
              % ','.join(disks))

    raidparts = []

    for disk in disks:
        if disk == disks[0]:
            part = 'part / --size=6000 --fstype=ext3 ' + \
                  '--ondisk=%s\n' % disk
            file.write(part)

            part = 'part /var --size=2000 --fstype=ext3 ' + \
                  '--ondisk=%s\n' % disk
            file.write(part)

        part = 'part raid.%s --size=5000 --ondisk=%s\n' % (disk, disk)
        file.write(part)

        raidparts.append('raid.%s' % disk)

    raid = 'raid /bigdisk --fstype=ext3 --device=md0 --level=1 %s\n' \
          % ' '.join(raidparts)
    file.write(raid)

```

```

def doDisk(file, disk):
    file.write('clearpart --all --initlabel --drives=%s\n' % disk)
    file.write('part / --size=6000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part /var --size=2000 --fstype=ext3 --ondisk=%s\n' % disk)
    file.write('part swap --size=2000 --ondisk=%s\n' % disk)
    file.write('part /mydata --size=1 --grow --fstype=ext3 --ondisk=%s\n'
              % disk)

#
# main
#
p = rocks_partition.RocksPartition()
disks = p.getDisks()

if nodename in [ 'compute-0-0' ]:
    file = open('/tmp/user_partition_info', 'w')

        if len(disks) == 2:
            doRaid(file, disks)
        elif len(disks) == 1:
            doDisk(file, disks[0])

    file.close()

```

5.5.3. Forcing the Default Partitioning Scheme for All Disks on a Compute Node

This procedure describes how to force all the disks connected to a compute node back to the default Rocks partitioning scheme regardless of the current state of the disk drive on the compute node.

The root disk will be partitioned as described in Default Partitioning and all remaining disk drives will have one partition with the name `/state/partition2`, `/state/partition3`, ...

For example, the following table describes the default partitioning for a compute node with 3 SCSI drives.

Table 5-2. A Compute Node with 3 SCSI Drives

Device Name	Mountpoint	Size
/dev/sda1	/	16 GB
/dev/sda2	swap	1 GB
/dev/sda3	/var	4 GB
/dev/sda4	/state/partition1	<i>remainder of root disk</i>
/dev/sdb1	/state/partition2	<i>size of disk</i>
/dev/sdc1	/state/partition3	<i>size of disk</i>

Create a new XML configuration file that will *replace* the current `custom-partition.xml` configuration file:

```
# cd /export/rocks/install/site-profiles/7.0/nodes/
```

```
# cp /export/rocks/install/rocks-dist/x86_64/build/nodes/custom-partition.xml replace-custom-part
```

Inside `replace-custom-partition.xml`, add the following section:

```
<pre>
echo "rocks force-default" > /tmp/user_partition_info
</pre>
```

Then apply this configuration to the distribution by executing:

```
# cd /export/rocks/install
# rocks create distro
```

To reformat compute node `compute-0-0` to your specification above, you'll need to first remove the partition info for `compute-0-0` from the database:

```
# rocks remove host partition compute-0-0
```

Then you'll need to remove the file `.rocks-release` from the first partition of *each disk* on the compute node. Here's an example script:

```
for file in $(mount | awk '{print $3}')
do
  if [ -f $file/.rocks-release ]
  then
    rm -f $file/.rocks-release
  fi
done
```

Save the above script as `/share/apps/nukeit.sh` and then execute:

```
# ssh compute-0-0 'sh /share/apps/nukeit.sh'
```

Then, reinstall the node:

```
# rocks set host boot action=install compute-0-0
# ssh compute-0-0 'shutdown -r now'
```

After you have returned all the compute nodes to the default partitioning scheme, then you'll want to remove `replace-custom-partition.xml` in order to allow Rocks to preserve all non-root partition data.

```
# rm /export/rocks/install/site-profiles/7.0/nodes/replace-custom-partition.xml
```

Then apply this update to the distribution by executing:

```
# cd /export/rocks/install
# rocks create distro
```

5.5.4. Forcing Manual Partitioning Scheme on a Compute Node

This procedure describes how to force a compute node to always display the manual partitioning screen during install. This is useful when you want full and explicit control over a node's partitioning.

Create a new XML configuration file that will *replace* the current `custom-partition.xml` configuration file:

```
# cd /export/rocks/install/site-profiles/7.0/nodes/
# cp /export/rocks/install/rocks-dist/x86_64/build/nodes/custom-partition.xml replace-custom-part
```

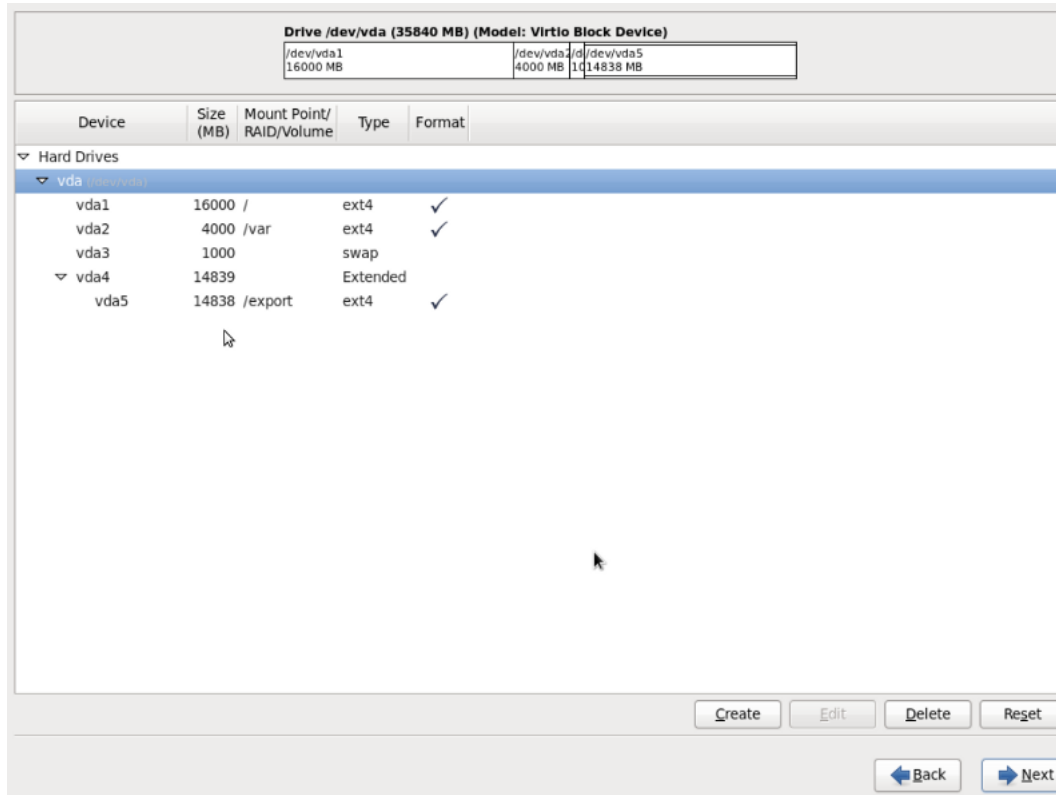
Inside `replace-custom-partition.xml`, add the following section:

```
<pre>
echo "rocks manual" > /tmp/user_partition_info
</pre>
```

Then apply this configuration to the distribution by executing:

```
# cd /export/rocks/install
# rocks create distro
```

The next time you install a compute node, you will see the screen:



To interact with the above screen, from the frontend execute the command:

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

5.6. Creating a Custom Kernel RPM

5.6.1. Creating a Custom Kernel RPM using kernel.org's Source

- On the frontend, check out the Rocks source code. See [Access to Rocks Source Code](#) for details.
- Change into the directory:

```
# cd rocks/src/roll/kernel/src/kernel.org
```

- Download the kernel source tarball from kernel.org. For example:

```
# wget http://www.kernel.org/pub/linux/kernel/v2.6/linux-2.6.24.4.tar.gz
```

- Create a kernel "config" file and put it in `config-<version>`

You can create the config file by using the following procedure:

```
# tar xzf linux-2.6.24.4.tar.gz
# cd linux-2.6.24.4
# make menuconfig
```

Configure the kernel anyway you need, and after the configuration is over choose to save the configuration in an alternative location. Enter the name of the file as `../config-2.6.24.4`. Finally, exit the configuration and remove the `linux-2.6.24.4` directory.



The `<version>` number must match the version number of the kernel source. For example, if you downloaded `linux-2.6.24.4.tar.gz`, the name of the config file must be `config-2.6.24.4`.

- Update `version.mk`.

The file `version.mk` has the following contents:

```
NAME           = kernel
RELEASE        = 1

VERSION       = 2.6.24.4
PAE            = 0
XEN            = 0
```

The `VERSION` value must match that of the linux kernel tarball you downloaded (e.g., 2.6.24.4).

If you are building a kernel for an i386 system that has more than 4 GB, you'll need to set the PAE (page address extension) flag. This will name the resulting kernel `kernel-PAE*rpm`. If the anaconda installer detects more than 4 GB memory, then it will install the kernel-PAE RPM and not the kernel RPM.

```
PAE            = 1
```

If you are building a kernel that contains Xen support, then set the XEN flag. This will name the resulting kernel `kernel-xen*rpm`.

```
XEN            = 1
```

It is illegal to set both the PAE and XEN flags.

If you want to build a kernel that the installer will use, then you must unset both the PAE and XEN flags (the default configuration of `version.mk`). This will name the resulting kernel `kernel*rpm`.



If you want to build a custom kernel for both the installer and for the running system and the running system requires either the kernel-PAE or kernel-xen RPM, then you'll have to execute this procedure twice: once to build the installer kernel (PAE = 0 and XEN = 0) and once to build the kernel that will be used on the running system (PAE = 1 or XEN = 1).

- Build the kernel:

```
# make rpm
```

- Copy the resulting RPMs into the current distribution:


```
# cp ../../RPMS/<arch>/kernel*rpm /export/rocks/install/contrib/7.0/<arch>/RPMS/
```

Where *<arch>* is i386 or x86_64.

- Rebuild the distribution:

```
# cd /export/rocks/install
# rocks create distro
```

- Test the new kernel by reinstalling a compute node:

```
# shoot-node compute-0-0
```

- If the kernel works to your satisfaction, reinstall all the compute nodes that you want to run the new kernel.

5.7. Enabling RSH on Compute Nodes

The default Rocks configuration does not enable rsh commands or login to compute nodes. Instead, Rocks uses ssh as a drop in replacement for rsh. There may be some circumstances where ssh does not have exactly the same semantics of rsh. Further, there may be some users that cannot modify their application to switch from rsh to ssh. If you are one of these users you may wish to enable rsh on your cluster.



Enabling rsh on your cluster has serious security implications. While it is true rsh is limited to the private-side network this does not mean it is as secure as ssh.

Enabling rsh is done by setting an *attribute*. To enable rsh on all compute nodes, execute:

```
# rocks set appliance attr compute rsh true
```

To apply this configuration change to the compute nodes, reinstall all your compute nodes.

If you only want to enable rsh on a specific node (e.g., compute-0-0), execute:

```
# rocks set host attr compute-0-0 rsh true
```

To apply this configuration change to compute-0-0, reinstall compute-0-0.

5.8. Hostbased vs. Key-base SSH Authentication

The default Rocks configuration uses host-based ssh key authentication. Prior to Rocks 6.1, supported only user-defined ssh keys. The advantage of host-based authentication is that a common file system for user home areas (that hold `.ssh/authorized_keys`) is no longer needed. Root still requires key-based authentication on all nodes.



Rocks defines a single SSH hostkey for all subordinate nodes. This can be overridden by the user to specify a unique hostkey for every single node. The frontend host key is different from the rest of the cluster.

Hostbased authentication is controlled by the contents of `/etc/ssh/shosts.equiv`. This file is served by the 411 service in Rocks. This is automatically updated when nodes are added or removed or when `rocks sync config` is run.

To turn off host-based SSH authentication, execute the following on your frontend

```
# rocks set attr rocks_autogen_user_keys true
# rocks sync config
```

this will generate an empty `/etc/ssh/ssh_host_key` file and publish via 411. It will also create the file `/etc/profile.d/rocks_autogen_user_keys`. The existence of that file is tested for in `/etc/profiles/ssh-keygen.sh`.

5.8.1. Per-host SSH Keys

SSH hostkeys are stored in the Rocks Secure Attributes database. By default, Rocks generates a single cluster-wide SSH key. The stored key is placed on the host after first boot via scp. To generate a key for a specific key for a particular node, say `compute-0-0`, repeat the following recipe for every node.

```
/usr/bin/ssh-keygen -q -t rsa -N "" -f /root/rsakey.tmp
/usr/bin/ssh-keygen -q -t dsa -N "" -f /root/dsakey.tmp
/opt/rocks/bin/rocks add host sec_attr compute-0-0 ssh_host_rsa_key value=/root/rsakey.tmp crypte
/opt/rocks/bin/rocks add host sec_attr compute-0-0 ssh_host_rsa_key.pub value=/root/rsakey.tmp.pu
/opt/rocks/bin/rocks add host sec_attr compute-0-0 ssh_host_dsa_key value=/root/dsakey.tmp crypte
/bin/rm /root/dsakey.tmp /root/rsakey.tmp
```

5.9. Adding a New Appliance Type to the Cluster

This procedure describes how to add a new appliance type to your cluster. This is useful when you want a subset of compute nodes to have specific behavior that is different from the rest of the compute nodes. For example, if you want all the nodes in cabinet 1 to be configured differently from the rest of the compute nodes.

Before you begin, you'll want to be comfortable with the Rocks XML framework that is used to produce a configuration graph. Details on this framework are found in the [Developers Guide](#)².

First, you'll need to create a new node XML file. This file will contain the configuration scripts and/or packages that will be applied to each of your appliances. Let's call it `my-compute.xml`. This file should be created in the directory `/export/rocks/install/site-profiles/7.0/nodes`. Below is the contents of the file:

```
<?xml version="1.0" standalone="no"?>

<kickstart>

<description>
My specialized compute node
</description>

<changelog>
</changelog>

<post>

<file name="/etc/motd" mode="append">
My Compute Appliance
</file>
```

```
</post>

</kickstart>
```

Now, we'll link the above file into the existing XML configuration graph. We'll simply point the above XML node to the existing `compute.xml` node. In object-oriented terms, we are inheriting all the functionality of the compute appliance and then extending it.

To link `my-compute.xml` to `compute.xml`, in the directory `/export/rocks/install/site-profiles/7.0/graphs/default`, create the file `my-appliance.xml` and have it contain:

```
<?xml version="1.0" standalone="no"?>

<graph>

<description>
</description>

<changelog>
</changelog>

<edge from="my-compute">
    <to>compute</to>
</edge>

<order gen="kgen" head="TAIL">
    <tail>my-compute</tail>
</order>

</graph>
```

To apply the changes above to the current distribution, execute:

```
# cd /export/rocks/install
# rocks create distro
```

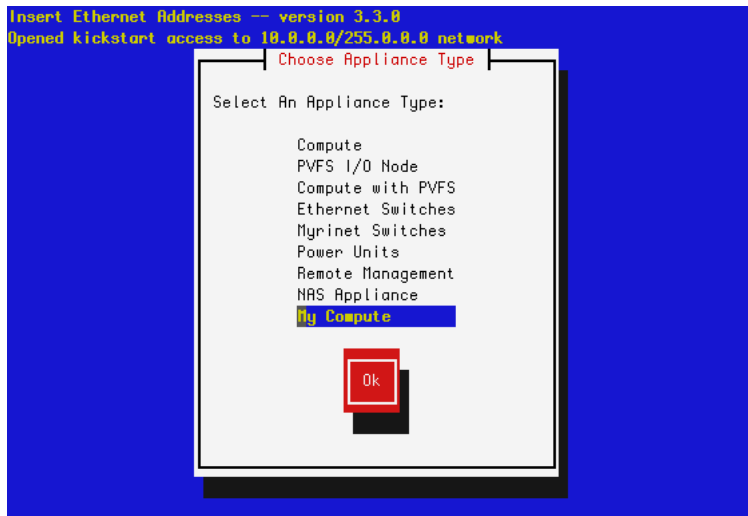
Now we need to add an entry into the Rocks MySQL database. This is accomplished with the `rocks` command line:

```
# /opt/rocks/bin/rocks add appliance my-compute membership='My Compute' \
    node='my-compute'
```

Now let's retarget an existing compute node. We'll use `insert-ethers` to accomplish this task. First, ask `insert-ethers` to replace `compute-0-0`:

```
# insert-ethers --replace compute-0-0
```

This displays the screen:



Select *My Compute* then hit *Ok*. This removes `compute-0-0` from the database and the next node that asks to be configured (that is, the next node that sends out a DHCP request) will be assigned the name `my-compute-0-0`. To see this in action, instruct `compute-0-0` to reinstall itself:

```
# shoot-node compute-0-0
```

Eventually, you'll see `insert-ethers` report that it discovered `my-compute-0-0`. After the node installs, it will be configured as a *my-appliance*. You can login to the node by executing:

```
# ssh my-compute-0-0
```

Your custom appliance can be applied to any new node in your system by starting `insert-ethers` as instructed above, then by booting a new node in configuration mode (by forcing it to PXE boot or by booting the node with the Kernel/Boot Roll).

5.10. Adding a Device Driver

This section describes how to add a device driver to the installation environment (`initrd.img-7.0-<arch>`). This enables the installation environment to use the new driver as well as installing the device driver into the running environment (that is, after the node has installed).

This feature is enabled by `ddiskit`³ which is maintained by Jon Masters at Red Hat.

1. Set up a build environment:

```
# cd /export
# mkdir src
# git clone http://git.rocksclusters.org/git/core/kernel/.git
# cd kernel
# git checkout -b ROCKS_<major_version>_<minor_version>
```

2. Go to the directory which holds the device driver code:

```
# cd src/rocks-boot/enterprise/5/images/drivers
```

3. In this directory, you'll see some example drivers. Let's look at the *e1000* driver:

```
# cd e1000
```

4. If you want to supply a new version of the driver, you'll have to download the e1000 source tarball and copy the *.c and *.h files from the tarball to this directory. Make sure all the *.c and *.h files are listed at the top of the *Makefile*:

```
MODULES := e1000

SOURCES := e1000_main.c e1000_hw.c e1000_ethtool.c e1000_param.c

HEADERS := e1000.h e1000_hw.h e1000_osdep.h
```

5. You'll need to make sure the proper PCI ids are in the file *pcitable*. For example, to test on one of our Dell SC1425's, we added the line:

```
0x8086 0x1076 "e1000" "Intel|82541GI/PI Gigabit Ethernet Controller (rev 05)"
```

6. Now we'll need to specify to the device driver building code that the e1000 driver should be built. To do this, edit the file *subdirs*:

```
# cd ..
# vi subdirs
```

7. Change the section from:

```
#
# put a list of all the driver directories that you'd like to build.
#
# for example, to build the 'e1000' driver, uncomment the line below:
#e1000
```

to:

```
#
# put a list of all the driver directories that you'd like to build.
#
# for example, to build the 'e1000' driver, uncomment the line below:
e1000
```

8. Build the *rocks-boot* package:

```
# cd /export/kernel/src/rocks-boot
# make rpm
```

9. When this completes, copy the binary RPMs into a directory where the distribution building utility (*rocks-dist*) will find and include them:

```
# cp /export/kernel/RPMS/x86_64/rocks-boot* \
/export/rocks/install/contrib/7.0/x86_64/RPMS/
```



If you are building on an i386 system, change the above x86_64 references to i386.

10. Rebuild the distro:

```
# cd /export/rocks/install
# rocks create distro
```

11. Install the newly created *initrd.img-7.0-<arch>* and its matching kernel *vmlinuz-7.0-<arch>* so PXE booted nodes will get the new device drivers:

```
# cd /export/rocks/install
# rpm -Uvh --force rocks-dist/x86_64/RedHat/RPMS/rocks-boot-5*.rpm
# cp /boot/kickstart/default/initrd.img-7.0-<arch> /tftpboot/pxelinux/
# cp /boot/kickstart/default/vmlinuz-7.0-<arch> /tftpboot/pxelinux/
```

12. Now PXE boot a node. This node will load your new driver and will install this driver into the running environment.

5.10.1. Adding a New Device Driver (That Isn't One of the Example Drivers)

If the name of your device driver you wish to add is not one of the example device drivers (e.g., `ata_piix`, `cciss`, `e1000`, `sk98lin`, or `tg3`), then you'll need to create a new directory and populate it with the appropriate files.

For help on how to populate your new driver directory, see:
`/export/kernel/src/rocks-boot/enterprise/5/images/drivers/INSTALL`.

The rest of the build process follows the same procedure as above starting at step 6 where you'll have to add the name of your new driver to the file `subdirs`.

5.11. Extending DNS

Rocks provides a mechanism to put external hostnames under the DNS control of your cluster. Generally, external hosts have names served by site-wide DNS servers. However if there is no external DNS server available, you may want to use your frontend's DNS server to handle the name to IP mappings for certain non-cluster nodes.

Since the DNS configuration file is automatically generated by Rocks, you cannot add static configuration to the standard zone files in `/var/named`. Instead, you need to put local name mappings in the file:

```
/var/named/<networkname>.domain.local
```

And reverse mappings (IP to name) in:

```
/var/named/reverse.<networkname>.domain.<networkprefix>.local
```

Where *networkname* for a default Rocks installation is *private* and *networkprefix* for a default Rocks installation is *1.10*. That is, the above file names would be:

For forward mapping: `/var/named/private.domain.local` For reverse mapping:
`/var/named/reverse.private.domain.1.10.local`

To add the mappings in the above `.local` files to the running configuration, execute:

```
# rocks sync config
```

These files are in the BIND configuration format, just like the standard `private.domain` and `reverse.private.domain` files.



Your external hosts will have names in the `.local` cluster domain.



Errors in your local DNS files will cause the entire local cluster domain naming to fail. Proceed with caution.

5.12. Changing the Root password

Starting with Rocks v5.4.3, we have attempted to improve the security of the Rocks system by making changes to the way some of the important passwords are generated and stored.

To change the root password on the frontend, you can now run the **passwd** command.

To change the root password on the client nodes, you can use the rocks secure attribute infrastructure.

Example 5-1. Change the root password of a single compute host

```
# rocks set host sec_attr compute-0-0 attr=root_pw
  Enter root_pw:
Confirm root_pw:
# rocks list host sec_attr compute
HOST      ATTR      VALUE                                     ENC
compute-0-0: root_pw $1$G55uZRrO$5MI7Nv14U5yWfKAlauqPE0 crypt
# rocks sync host sec_attr compute-0-0
```

For more information about changes to the rocks security infrastructure, refer to the Rocks Password Infrastructure section of this usersguide



The default rocks password policy is to set the root password of the frontend to be the password supplied during the installation. For client nodes, the root password will be generated at random during installation. Unless there is compelling reason to do so, (such as giving a user root access to a compute node) the Rocks group advises against changing the root password on the client nodes.

5.13. The insert-ethers command line

This section describes the various command line argument accepted by insert-ethers which can be used to customize hostname, IPs address and other parameters of the discovered node during installation.

- insert-ethers

```
Insert Ethernet Addresses -
Usage: insert-ethers [-hv] [-p password] [-u host] [-d database] [--help]
  [--list-rcfiles] [--list-project-info] [--verbose] [--update] [--staticip]
  [--rcfile arg] [--host host] [--password password] [--db database]
  [--user host] [--baseip ip address] [--basename basename] [--hostname hostname]
  [--ipaddr ip address] [--cabinet number] [--rack number] [--inc number]
  [--rank number] [--replace hostname] [--remove hostname]
  [--membership membership] [--os the OS to install on the machines]
```

- *[-p password] [-u host] [-d database]*: they are used to change the hostname the password and the name of the backend database used by insert-ether and Rocks. Users should not use this arguments.
- *[--list-rcfiles]*: it lists the RC file used by insert-ethers (RC file are internal configuration file used by Rocks)
- *[--list-project-info]* it returns the Rocks version and the Roll this command belongs to
- *[--baseip ip address]* the IP address that will be used for the first node, when installing multiple nodes insert-ethers start from this IP address and then it uses --inc to get the following IP addresses. Default is the highest IP address available in the private network. If the private network is 10.1.0.0 (with netmask 255.255.0.0) it will start from 10.1.255.255.

- *[--inc number]* the increment value to use when determining the following IP address during the installation of multiple nodes. Default to -1. So if the start IP address is 10.1.255.255 the following IP address will be 10.1.255.254, 10.1.255.253, 10.1.255.252, and so on.
- *[--basename basename]* this is the first part of the final hostname of the hosts that will be installed. The final hostname will be in the form of <basename>-<rack>-<rank> where rack and rank are two counters internally calculated by rocks.
- *[--cabinet number] [--rack number]* these two options are equivalent. They can be used to increment the 'rack' part of the final hostname (hostname are in the form of <basename>-<rack>-<rank>). `insert-ethers --cabinet 2` selecting compute node will install nodes with the hostname equal to: `node-2-0`, `node-2-1`, `node-2-2`, etc.
- *[--rank number]* this can be used to set the rank part of the hostname of the installed node(). This number default to 0 at the first installed node and then it is automatically incremented by Rocks, but if the user wants to skip some number or wants to start from a number different from 0, it can use this option.
- *[--hostname hostname]* this is the full hostname that will be assigned to the first host you will install with `insert-ethers`. If you specify this option `insert-ethers` will install only one host.
- *[--ipaddr ip address]* it specifies the IP address that will be assigned to the first installed host. If you specify this option `insert-ethers` will install only one host.
- *[--replace hostname]* it deletes hostname from the Rocks database and install 1 new discovered host with the same rack and rank of the deleted host. This command will install only one host. `return`.
- *[--remove hostname]* it removes hostname from the rocks database this command is equivalent to 'rocks remove host hostname'. This command will not start the GUI, it will simply delete the hostname and return.
- *[--membership membership]* You can specify the membership name and skip the first prompt screen.
- *[--os the OS to install on the machines]* it specifies the OS to install on the new discovered host. Default is linux.

Notes

1. http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Installation_Guide/s1-kickstart2-options.html#s2-kickstart2-options-part-example
2. <http://www.rocksclusters.org/roll-documentation/developers-guide/>
3. <http://dup.et.redhat.com/ddiskit/>

Chapter 6. Community Resources

6.1. Access to Rocks Source Code

The source code for Rocks is available from a public read-only GIT¹ repository. As of Rocks 5.7 GIT (git) is installed on your frontend machine, for older version of Rocks (or non-Rocks platforms) you will need to install the Git package yourself. You can download rpm packages from the EPEL Red Hat web site at http://dl.fedoraproject.org/pub/epel/5/x86_64/repoview/git.html

Git is a distributed source control system. Several tutorials are provided on the GIT web site <http://git-scm.com/documentation>. Fundamentally, distributed source control turns your checked out source code from a read-only copy into a repository and then grants read-write access to others based on your checked out repository. In the context of Rocks this means individual Roll developers can independently control who has access to their code.

You can use the Github Web interface to view the source code using your browser at <https://github.com/rocksclusters/>. Each Roll has its own repository named after the Roll. So if you need to checkout the Base roll you should checkout the repository called base under the Rocks Github web page (in this example <https://github.com/rocksclusters/base>) To check out the current source code you will first need to clone one of the read-only repository listed on the main web page of the Github web site. This operation can take several minutes to complete.

```
$ git clone http://github.com/rocksclusters/base;
Cloning into 'base'...
$ cd base
$ git status
# On branch master
nothing to commit (working directory clean)
$ vi DESCRIPTION
$ git status
# On branch master
# Changes not staged for commit:
#   (use "git add <file>..." to update what will be committed)
#   (use "git checkout -- <file>..." to discard changes in working directory)
#
# modified:   DESCRIPTION
#
no changes added to commit (use "git add" and/or "git commit -a")
```

6.2. All Past Rocks Releases

All past Rocks releases can be found at <ftp.rocksclusters.org:/pub/rocks.6>

6.3. Email Discussion List

The primary mode for support for Rocks is through our email discussion list.

Click [here](#)⁷ to subscribe.

The archives⁸ hold messages from the Rocks discussion list dating back to June 2002.

If you need to search old messages you can use the following links:

- Mailing list ARChives⁹
- Google¹⁰ with the extra directive `site://lists.sdsc.edu/pipermail/npaci-rocks-discussion` in the search bar (to limit search results to Rocks mailing list only).

Notes

1. <http://git-scm.com/>
2. http://dl.fedoraproject.org/pub/epel/5/x86_64/repoview/git.html
3. <http://git-scm.com/documentation>
4. <https://github.com/rocksclusters/>
5. <https://github.com/rocksclusters/base>
6. <ftp://ftp.rocksclusters.org/pub/rocks>
7. <https://lists.sdsc.edu/mailman/listinfo/npaci-rocks-discussion>
8. <https://lists.sdsc.edu/pipermail/npaci-rocks-discussion/>
9. <http://marc.info/?l=npaci-rocks-discussion&r=1&w=2>
10. <https://www.google.com/search?q=site%3Ahttps%3A%2F%2Flists.sdsc.edu%2Fpipermail%2Fnpaci-rocks-discussion>

Chapter 7. Administration Examples

7.1. Introduction to the Rocks Command Line

In Rocks 4.3 the Rocks command line was introduced to provide a more uniform interface to the underlying structures used to control system configuration and behaviour. Wherever possible, Rocks uses a SQL database (MySQL currently) to hold information about nodes, partitioning information, boot parameters and other information. Based on information in the database, various configuration files are rewritten. The re-generation of configuration files occurs everytime a node is added or deleted from the cluster. The re-generation of configuration files can also be forced. A large fraction of Rocks commands manipulate data held in the configuration database. In general, the process of changing configuration is a two-step process:

1. Use rocks commands to change configuration in the database (e.g. `rocks set host`)
2. Rewrite configuration files using `rocks sync config`

It should be noted that step 1 above is usually called several times to update in the database and then step 2 is called to write individual configuration files in the format that the native OS tools understand.



Rocks commands have arguments and parameters. Parameters are of the form "param=<value>" and may appear anywhere. Arguments must appear in the order defined by the command. To get help on any rocks command type "help" for the argument to the command. For example `rocks set host interface ip help`

7.2. Boot Order and PXE First

Prior to Rocks 4.3, the BIOS-defined boot order of a compute node *required* that a network boot (known as PXE) come after local hard disk. In particular the boot order in BIOS would be set as

1. CDROM
2. Hard Disk
3. On-board Network Device (PXE)

A user would have to intercept the boot sequence (often by hitting the F12 key on an attached keyboard) to force a network boot. Rocks also provided a small utility on each node (`/boot/kickstart/cluster-kickstart-pxe`) that would manipulate the two-bytes on the local hard disk to force BIOS to bypass booting from the local disk and try the next device on the boot list. When the boot order was set as above, the node would PXE boot and therefore re-install.

The logic for this structure was that a frontend did not need to know the state of node (whether it had failed and should be reinstalled or had some other intermediate state). Also it is not required that a frontend be up for a node to reboot itself. Another practical issue arises for PXE booting large clusters. Since the PXE client is in NIC firmware, no assumptions about timeouts, retries or other elements that figure into robustness could be made. Large cluster reinstalls (or reboots) for a kernel that comes over PXE would often result in hung nodes because of the low level of robustness of TFTP (the underlying protocol used to transfer initial kernel and ramdisk image for nodes booting over the network). For wholesale re-installation of large clusters, PXE does not scale well. For this, Rocks provides the installation kernel and initial ramdisk image on the local hard drive. The

command `/boot/kickstart/cluster-kickstart` run on a local node will cause that node to re-install itself by using a local (hard disk) copy of the installation kernel and initial ramdisk.



The above boot order and behaviour continues to be supported in Rocks 4.3. That is, existing rocks clusters can be upgraded without requiring the cluster owner to change any BIOS settings.

7.3. Support for PXE First

Rocks supports a network device first (or PXE first) BIOS-defined boot order. It is now *recommended* that a network boot (known as PXE) come before local hard disk. In particular the boot order in BIOS should be set as

1. CDROM
2. On-board Network Device (PXE)
3. Hard Disk

The default PXE "action" is to simply pass to the next device down on the BIOS boot list. In the usual case, this is to the local hard disk. Most of the time decision to boot or reinstall is still left to the local node and the frontend does not need to know which state the node desires. If booting into re-installation (e.g. the node either did not shut down properly, or `/boot/kickstart/cluster-kickstart` was called locally) that will proceed as expected. However, it is possible to change this action on a per-node basis.

7.4. Forcing a Re-install at Next PXE Boot

Starting with Rocks 4.3, the frontend must be configured to tell a node to re-install at the next PXE boot. This action is controllable on a per-node basis. At the end of successful installation, the node requests the frontend to set its PXE boot to *os*. To re-install a node using PXE (e.g. `compute-0-0`), then do the following:

```
# rocks set host boot compute-0-0 action=install
# ssh compute-0-0 "shutdown -r now"
```



If the boot order has not been set to PXE first, you can force a PXE boot with the local keyboard, or by calling `/boot/kickstart/cluster-kickstart-pxe` on the local node.

7.5. Inspecting and Changing PXE Behaviour

There are two parts to the Rocks database for modifying PXE behaviour: *boot* and *bootaction*. The "boot" part determines which logical action should be performed. The two common actions are "os" and "install." The second table is the bootaction table. This associates a logical action with a specific TFTP configuration.



It is possible to have commands affect all nodes. In this case use '%' as the host wildcard. For example `rocks set host boot % action=install` will cause ALL nodes to reinstall the next time they boot.



For commands that take lists of hosts, it is possible to use an appliance type for the host(s) argument. `rocks list appliance` are the list of valid appliance types. To set the boot action of all compute appliances to be *install*, use `rocks set host boot compute action=install`.

The following illustrates how to inspect the current action of nodes and then the specifics of each action.

```
# rocks list host boot
HOST          ACTION
vizzy:       os
compute-0-0: os
compute-0-1: os
compute-1-0: os
compute-1-1: install
compute-2-0: os
compute-2-1: os

# rocks list bootaction output-col=action,kernel
ACTION          KERNEL
install:        vmlinuz-5.2-i386
install headless: vmlinuz-5.2-i386
memtest:        kernel memtest
os:             localboot 0
pxeflash:       kernel memdisk bigraw
rescue:         vmlinuz-5.2-i386
```

In the above, all nodes are set to boot the "os", except for node `compute-1-1`. That node will call the boot action named "install". In the case the TFTP configuration file contain the details arguments of the listed in the install action. The command `rocks list bootaction` shows the details of each logical action.

7.5.1. Changing a logical PXE action

It is possible to override the details of a logical action on a per-node basis. Suppose that we wanted to make the logical action of "install" for `compute-1-1` to be headless and to set a flag `acpi=off`. Then the following will accomplish this:

```
# rocks add bootaction action="install headless noacpi" kernel="vmlinuz-5.2-i386" ramdisk="initrd
args="ks ramdisk_size=150000 lang= devfs=nomount pxe kssendmac selinux=0 noipv6 headless vnc acpi

# rocks set host installaction compute-1-1 action="install headless noacpi"
```

To inspect that the change is indeed specific to just `compute-1-1`, then do the following

```
# rocks list host compute-1-1
MEMBERSHIP  CPUS  RACK  RANK  RUNACTION  INSTALLACTION
Compute     2    0    0    os         install headless noacpi

# rocks list host compute-1-0
```

```
MEMBERSHIP  CPUS  RACK  RANK  RUNACTION  INSTALLACTION
Compute     8     0     1     os         install
```

In the above, compute-1-1 has a specific override for its install action, where compute-1-0 still retains the default install action.

7.5.2. Running Memtest86

It is often useful to run the memory testing tool `memtest86+`¹ to determine if memory is valid. The straightforward way to accomplish this is to apply the following procedure (in our example case for host compute-1-1)

1. `# rocks set host runaction compute-1-1 action=memtest`
`# rocks set host boot compute-1-1 action=os`
2. Boot node compute-1-1 by power cycle or other means.
3. After compute-1-1 has successfully started the diagnostic, reset the runaction parameter:
`# rocks set host runaction compute-1-1 action=os`

7.6. Working with and Modifying Network Configuration

The Rocks database holds information that has been discovered about a host and in particular records network interface information including MAC addresses and local device modules. The Rocks command line has several tools to inspect and update entries in the database.

For the following discussion, a compute node with a public interface will be used for illustration. To list the ethernet interfaces do the following:

```
# rocks list host interface compute-0-0

SUBNET  IFACE  MAC                IP                NETMASK          MODULE  NAME                VLAN  O
private eth0   00:0e:0c:a7:57:d7  10.1.255.254     255.255.0.0     -----  compute-0-0        ----  -
public  eth1   00:19:b9:21:b8:b6  137.110.119.93  255.255.255.0   -----  kong.rocksclusters.org  ----  -
```

Suppose that it is desired to swap these two interfaces. That is, we would want to associate the private network with eth1 and the public network with eth0. The following will change the information only in the database.

```
# rocks swap host interface compute-0-0 ifaces=eth0,eth1

# rocks list host interface compute-0-0
SUBNET  IFACE  MAC                IP                NETMASK          MODULE  NAME                VLAN  O
public  eth0   00:0e:0c:a7:57:d7  137.110.119.93  255.255.255.0   -----  kong.rocksclusters.org  ----  -
private eth1   00:19:b9:21:b8:b6  10.1.255.254     255.255.0.0     -----  compute-0-0        ----  -
```

We see that eth0 is now associated with the public network, and eth1 is associated with the private network.

To apply the changes to compute-0-0, execute:

```
# rocks sync config
# rocks sync host network compute-0-0
```

7.7. Reinstall All Compute Nodes with SGE

This section describes how to reinstall all the nodes under the control of SGE.

As root, execute:

```
# /opt/gridengine/examples/jobs/sge-reinstall.sh
```

This will produce output like:

```
Your job 1 ("reboot.qsub") has been submitted
Set compute-0-0.local for Reinstallation
Your job 2 ("reboot.qsub") has been submitted
Set compute-0-1.local for Reinstallation
Your job 3 ("reboot.qsub") has been submitted
Set compute-0-2.local for Reinstallation
```

For each compute node, the script submits a high-priority reinstall job. The reinstall jobs will move ahead of the currently queued jobs (running jobs will not be disturbed). When a busy compute node becomes free, the next job submitted to it will be the reinstall job. All remaining jobs will remain queued until the node completes its installation, after which SGE will recognize the node as up and then submit a previously queued job to it.

Notes

1. <http://www.memtest.org>

Chapter 8. Advanced Tasks

8.1. Managing the Firewall on the Cluster Nodes

The firewall for all cluster nodes (including the frontend) is managed with the Rocks command line. As of Rocks 5.4.3, all firewall rules must have a name associated. The name is used as the handle for deleting and redefined rules. For example, a rule named 'MYRULE' can be defined globally, but a particular host can redefine 'MYRULE' to do something else. As of Rocks 5.4.3 rules are associated by category and single 'add firewall' command handles global, os, appliance, and host categories.

8.1.1. Categories. Subtle Syntax

In Rocks 5.4.3 and newer releases, the concept of "categories" is used by the database. Prior to Rocks 5.4.3, there were several different add firewall commands. Namely,

1. rocks add firewall
2. rocks add os firewall
3. rocks add appliance firewall
4. rocks add host firewall

Logically, these were defined to add a firewall command to one of [global, os , appliance, host]. When firewall rules were resolved from the database, rules would be inherited from global, then os, then appliance, and finally host-specific rules.



As of 5.4.3, these have been replaced by a single command that looks like

```
# rocks add firewall category=index ...
```

Here categories are [global, os, appliance, host]. This means that adding firewall rule to a particular host (say, compute-0-0) becomes

```
# rocks add firewall host=compute-0-0 ...
```

This rule applies only to the compute-0-0 host. To add a rule to the Compute appliance becomes:

```
# rocks add firewall appliance=Compute ...
```

The appliance rule applies to all Compute appliances.

8.1.1.1. Listing Firewall Rules. Resolved Rules.

There are two commands for listing firewall rules for a host and they differ in subtle, but important, ways. Again we'll use compute-0-0 as an example.


```
# rocks list firewall host=compute-0-0
# rocks list host firewall compute-0-0
```

The first command lists all firewall rules that are specific to host compute-0-0 only. By default, this is the empty set. The second command list all *inherited* commands and is, by default not empty. For example,

```
# rocks list firewall host=compute-0-0
#
```

and

```
# rocks list host firewall compute-0-0 maxwidth=15
# rocks list host firewall compute-0-0-0 maxwidth=15
RULENAME          SERVICE PROTOCOL CHAIN  ACTION NETWORK OUTPUT-NETWORK FLAGS           COMMENT
A20-ALL-PRIVATE   all     all     ACCEPT INPUT  private -----
A20-SSH-PUBLIC     ssh     tcp     ACCEPT INPUT  public  ----- -m state --stat -----
A30-RELATED-PUBLIC all     all     ACCEPT INPUT  public  ----- -m state --stat -----
R900-PRIVILEGED-TCP all     tcp     REJECT INPUT  public  ----- --dport 0:1023 -----
R900-PRIVILEGED-UDP all     udp     REJECT INPUT  public  ----- --dport 0:1023 -----
```

Notice that in the second form, the source of the firewall (global, in this example) is listed so that an administrator can see where a rule is defined. One should also notice that list host firewall will output all possible firewall rules that *might* apply. For example, the above has rules defined for both the "public" and "private" networks in the global scope. If compute-0-0 does not have an interface on the public network, then those rules will not appear in the final firewall configuration.

The list commands (like all list commands in Rocks) are intended for human readability. While the report commands are machine readable. The actual firewall rules written for compute-0-0 would be generated with `rocks report host firewall`. The report command resolves that actual interface for the named network. For example:

```
#rocks report host firewall compute-0-0
<file name="/etc/sysconfig/iptables" perms="500">
*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -i lo -j ACCEPT
# A20-ALL-PRIVATE (global) :
-A INPUT -i eth0 -j ACCEPT
COMMIT
</file>
```

A comment is put into this file for each applicable firewall rule. Notice that compute-0-0 has only a private network, so that all firewall rules bound to a public network are ignored. `rocks report host firewall` is used internally during node build and by `rocks sync host firewall`

8.1.2. Adding Rules. Ordering Rules. Examples

Firewall concepts are taken from linux iptables. The defined chains are INPUT, OUTPUT, FORWARD, ACCEPT, REJECT, DROP. The most common is to define how to match a packet on the INPUT chain and then either ACCEPT, REJECT, or DROP.

Overall order of rule definition matters and this is where the rulename is important. Suppose that on the frontend, we want to accept ftp connections from all hosts, instead of the rocks default. First lets list the rules for a frontend (COMMENT field has been edited out)

RULENAME	SERVICE	PROTOCOL	CHAIN	ACTION	NETWORK	OUTPUT-NETWORK	FLAGS
A10-REJECT-411-TCP	all	tcp	REJECT	INPUT	private	-----	--dport 372 -
A10-REJECT-411-UDP	all	udp	REJECT	INPUT	private	-----	--dport 372 -
A20-ALL-PRIVATE	all	all	ACCEPT	INPUT	private	-----	-----
A20-SSH-PUBLIC	ssh	tcp	ACCEPT	INPUT	public	-----	-m state --st
A30-RELATED-PUBLIC	all	all	ACCEPT	INPUT	public	-----	-m state --st
A40-HTTPS-PUBLIC-LAN	https	tcp	ACCEPT	INPUT	public	-----	-m state --st
A40-WWW-PUBLIC_LAN	www	tcp	ACCEPT	INPUT	public	-----	-m state --st
A50-FORWARD-RELATED	all	all	ACCEPT	FORWARD	public	private	-m state --st
A60-FORWARD	all	all	ACCEPT	FORWARD	private	-----	-----
MASQUERADE	nat	all	MASQUERADE	POSTROUTING	-----	public	-----
R10-GANGLIA-UDP	8649	udp	REJECT	INPUT	-----	-----	-----
R20-MYSQL-TCP	3306	tcp	REJECT	INPUT	-----	-----	-----
R30-FOUNDATION-MYSQL	40000	tcp	REJECT	INPUT	-----	-----	-----
R900-PRIVILEGED-TCP	all	tcp	REJECT	INPUT	public	-----	--dport 0:102
R900-PRIVILEGED-UDP	all	udp	REJECT	INPUT	public	-----	--dport 0:102

The R900-PRIVILEGED* rules are the last ones interpreted by iptables. They instruct iptables to reject all packets destined for privileged ports on the public network (--dport 0:1023). To accept ftp traffic (port 21) we need to add rules for both UDP and TCP traffic that are named *alphabetically* before the R900 rules. It is the rulename that determines order. Typically ACCEPT rules are labeled A<nn>- and REJECT rules are labeled R<nn>-. But that is only convention for Rocks system rules. Any name is valid.

```
# rocks add firewall host=frontend network=public protocol=tcp service=ftp chain=INPUT action=ACCEPT
# rocks add firewall host=frontend network=public protocol=udp service=ftp chain=INPUT action=ACCEPT
```



It is sometimes simpler on the screen to look at the actual machine report via `rocks report host firewall`

```
# rocks report host firewall frontend
<file name="/etc/sysconfig/iptables" perms="500">
*nat
# MASQUERADE (host) :
-A POSTROUTING -o eth1 -j MASQUERADE
COMMIT

*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -i lo -j ACCEPT
# A10-REJECT-411-TCP (host) :
-A INPUT -i eth0 -p tcp --dport 372 --sport 1024:65535 -j REJECT
# A10-REJECT-411-UDP (host) :
```

```

-A INPUT -i eth0 -p udp --dport 372 --sport 1024:65535 -j REJECT
# A100-PUBLIC-FTP-UDP (host) :
-A INPUT -i eth1 -p udp --dport ftp -j ACCEPT
# A20-ALL-PRIVATE (global) :
-A INPUT -i eth0 -j ACCEPT
# A20-SSH-PUBLIC (global) :
-A INPUT -i eth1 -p tcp --dport ssh -m state --state NEW -j ACCEPT
# A30-RELATED-PUBLIC (global) :
-A INPUT -i eth1 -m state --state RELATED,ESTABLISHED -j ACCEPT
# A40-HTTPS-PUBLIC-LAN (host) :
-A INPUT -i eth1 -p tcp --dport https -m state --state NEW --source &Kickstart_PublicNetwork;/&Ki
# A40-WWW-PUBLIC_LAN (host) :
-A INPUT -i eth1 -p tcp --dport www -m state --state NEW --source &Kickstart_PublicNetwork;/&Kick
# A50-FORWARD-RELATED (host) :
-A FORWARD -i eth1 -o eth0 -m state --state RELATED,ESTABLISHED -j ACCEPT
# A60-FORWARD (host) :
-A FORWARD -i eth0 -j ACCEPT
# A90-PUBLIC-FTP-TCP (host) :
-A INPUT -i eth1 -p tcp --dport ftp -j ACCEPT
# R10-GANGLIA-UDP (host) : block ganglia traffic from non-private interfaces
-A INPUT -p udp --dport 8649 -j REJECT
# R20-MYSQL-TCP (host) : block mysql traffic from non-private interfaces
-A INPUT -p tcp --dport 3306 -j REJECT
# R30-FOUNDATION-MYSQL (host) : block foundation mysql traffic from non-private interfaces
-A INPUT -p tcp --dport 40000 -j REJECT
# R900-PRIVILEGED-TCP (global) :
-A INPUT -i eth1 -p tcp --dport 0:1023 -j REJECT
# R900-PRIVILEGED-UDP (global) :
-A INPUT -i eth1 -p udp --dport 0:1023 -j REJECT

```



There are four things to notice. First, the ordering is alphabetical not numerical. A100 is after A10-R but before A20. A90 is after A60. Second, both of the accept rules for ftp come before the reject rules. Third, the public network has been resolved for this host to be eth1. See `rocks list host interface frontend`. And, fourth, each lists where it comes from.

Now apply the configuration to the host:

```
# rocks sync host firewall frontend
```

The host will now accept ftp traffic on its public interface.

8.1.2.1. How to affect a particular appliance

Suppose we wanted to close ftp on *all* networks for compute appliances (even private). First, let's look at the default rules for a compute node.

```

rocks list host firewall compute-0-0
RULENAME          SERVICE PROTOCOL CHAIN ACTION NETWORK OUTPUT-NETWORK FLAGS
A20-ALL-PRIVATE   all     all     ACCEPT INPUT  private -----
A20-SSH-PUBLIC     ssh     tcp     ACCEPT INPUT  public ----- -m state --state NEW
A30-RELATED-PUBLIC all     all     ACCEPT INPUT  public ----- -m state --state RELATE
R900-PRIVILEGED-TCP all     tcp     REJECT INPUT  public ----- --dport 0:1023
R900-PRIVILEGED-UDP all     udp     REJECT INPUT  public ----- --dport 0:1023

```

Here you can use the "appliance=<appliance-name>" argument for adding a firewall. Use `rocks list appliance` to get the valid names. Notice that by default all traffic is accepted on the compute node on the private network. If we want our REJECT rule to properly apply to all network interfaces, it must come before the A20-ALL-PRIVATE rule.

```
# rocks add firewall appliance=compute protocol=tcp service=ftp network=all chain=INPUT action=REJECT
# rocks add firewall appliance=compute protocol=udp service=ftp network=all chain=INPUT action=REJECT
```

This now rejects. FTP traffic. You can check the actual firewall rule for a particular compute host.

```
rocks report host firewall compute-0-0
<file name="/etc/sysconfig/iptables" perms="500">
*filter
:INPUT ACCEPT [0:0]
:FORWARD DROP [0:0]
:OUTPUT ACCEPT [0:0]
-A INPUT -i lo -j ACCEPT
# A19-REJECT-FTP-TCP (appliance) :
-A INPUT -p tcp --dport ftp -j REJECT
# A19-REJECT-FTP-UDP (appliance) :
-A INPUT -p udp --dport ftp -j REJECT
# A20-ALL-PRIVATE (global) :
-A INPUT -i eth0 -j ACCEPT
COMMIT
</file>
```

Now apply the firewall on all compute nodes.

```
# rocks sync host firewall compute
```

Firewall is now set to reject FTP on all networks for all compute nodes.

8.2. Enabling Two-Factor Authentication with Google Authenticator


Two-factor authentication is a technique to improve password security on frontends and login nodes. The two factors are something you know (password) and something you possess physically (e.g., a phone, a hardware key token, ...). Rocks supports two-factor authentication on Frontend and Login appliances using Google's free and open `google authenticator`¹. Two-factor authentication is often called one-time passwords (OTP), because the rotating codes (every 30 seconds, by default) are used just once and then discarded. This document uses two-factor and one-time passwords (OTP) interchangeably.

It is reasonable to ask, "How is this different from RSA or other hardware tokens?". In technique, this method is not different from hardware tokens. Both use time-based keys. However, standalone key tokens are arguably stronger because the device that you possess is not connected to any network (unlike your phone). Nevertheless, Google Authenticator is significantly stronger than just using password-only authentication.


8.2.1. Download and Install App for your Smartphone

Google authenticator requires the user to download a small application for their smartphone. There are apps for Android phones² and Apple iPhones³. Please see the appropriate link for your smartphone.

The Android Application is available at the Google Play Store



The IOS Application is available at the Apple iTunes Store



8.2.2. Creating a User's Second Factor (One Time Password)

Users must be in the `google-otp` group to utilize one-time passwords (OTP). The Rocks-defined PAM configuration first checks if the user is in the group, and then checks if a user-specific OTP configuration file is in the correct location (`/export/google-authenticator/<username>`). This enables cluster owners to have mixtures of users, some of whom use two-factor authentication and some of whom do not. It is highly recommended that users who can gain root privilege (`sudo` or direct login) utilize two-factor authentication.

The user creates a password file using `google-authenticator`. A number of options are asked of the user and it is recommended that the user selects "yes" to every question. The program will create a QR⁴ code that can be scanned by either the Android or IOS application to register the created password with the phone application.

```
% google-authenticator
```

The following example shows a screenshot of google-authenticator

```
File Edit View Terminal Help
```

```
[root@centos-6-3 google-otp]# google-authenticator
```

```
Do you want authentication tokens to be time-based (y/n) y
```



```
Your new secret key is: EETMAG56FNXY6PMY
```

```
Your verification code is 181697
```

```
Your emergency scratch codes are:
```

```
11781008
```

```
96354852
```

```
64212991
```

```
21210905
```

```
72522409
```

```
Do you want me to update your "/root/.google_authenticator" file (y/n) y
```


File Edit View Terminal Help



Your new secret key is: EETMAG56FNXY6PMY

Your verification code is 181697

Your emergency scratch codes are:

11781008

96354852

64212991

21210905

72522409

ADDITIONAL INFORMATION

Do you want me to update your `"/root/.google_authenticator"` file (y/n) y

code is a C library for encoding data in a QR Code symbol, a kind of 2D symbology that can be scanned by a camera or a QR code reader. It can store up to 7000 digits or 4000 characters.
Do you want to disallow multiple uses of the same authentication token? This restricts you to one login about every 30s, but it increases your chances to notice or even prevent man-in-the-middle attacks (y/n) y

By default, tokens are good for 30 seconds and in order to compensate for possible time-skew between the client and the server, we allow an extra token before and after the current time. If you experience problems with poor time synchronization, you can increase the window from its default size of 1:30min to about 4min. Do you want to do so (y/n) y


If the computer that you are logging into isn't hardened against brute-force login attempts, you can enable rate-limiting for the authentication module. By default, this limits attackers to no more than 3 login attempts every 30s. Do you want to enable rate-limiting (y/n) y

[root@centos-6-3 google-otp]#



Your terminal font may need to be minimized to fit the QR code onto your display so that it can be scanned by the Phone App

Scan the QR code with your Phone's Camera while in the Google Authenticator App. For example, in the iPhone app, you should get to screen that looks like



8.2.3. Enabling Two-factor for a particular user

Once a user has created a two-factor code (which is a private password file), it needs to be moved from the user's home directory to a more secure location. That password file, then needs to be owned by root for two-factor to become active. Removing this file, will disable two-factor authentication for a particular user.

```
# mv /home/<user>/.google-authenticator /export/google-authenticator/<user>
# chown root.root /export/google-authenticator/<user>
# make -C /var/411 clean groups
# rocks sync users
```



If the user has generated their password on a login node, then securely copy their `.google-authenticator` file to the `/export/google-authenticator` directory on the frontend. Rocks 411 will properly sync these files to all login appliances

8.2.4. The User Experience

If all has been configured correctly, then a user will see an extra step at login, a request for a verification code

```
[phil@build-x86-64 ~]$ ssh 192.168.122.12 -l phil
Verification code:
Password:
[phil@centos-6-3 ~]$
```

8.2.5. Frequently Asked Questions

1. Q: is root's OTP key file synced to Login node

A: No. Like other security keys associated with root on the frontend, this data is never transferred

2. Q: If I delete an OTP key file from the frontend and then run `rocks sync users`, will the file be deleted on the Login node?

A: Yes. The 411 handler on the Login node first deletes all keys (except root) in the `/export/google-authenticator` directory and then unpacks the new ones.

3. Q: Is it possible to turn off the automatic addition of users to the `google-otp` group?

A: Yes. This automatic addition is handled by a `rocks sync users` plugin. If you set the attribute `Info_GoogleOTPUsers` to "no", then no automated additions will be made. This will not remove users from the `google-otp` group that have already been added.

```
# rocks set attr Info_GoogleOTPUsers no
```

8.3. Flashing BIOS on Compute Nodes Using PXE

This procedure describes how to flash BIOS on a client node (e.g., compute node or viz tile node) by using the Rocks command line and PXE.

- First, you must add the BIOS flashing files (e.g., a DOS-based flash utility, .bin file or an autoexec.bat) to the directory `/opt/pxeflash/addon`.

As an example, to flash a Dell Dimension E521 desktop, we executed:

```
# cd /opt/pxeflash/addon
# wget http://ftp.us.dell.com/bios/DME521-010111.EXE
```

- To add the BIOS flashing file to a bootable floppy image, execute:

```
# cd /opt/pxeflash
# make build
```

This will incorporate all files in the `/opt/pxeflash/addon` directory into a floppy image named `pxeflash-FDSTD.288`.

- Now copy the floppy image into `/tftpboot/pxelinux/pxeflash`, by executing:

```
# make install
```

- Set a client node to PXE boot your flash image:

```
# rocks set host runaction <hostname> action=pxeflash
# rocks set host boot <hostname> action=run
```

- PXE boot the client node. After the client boots, it will display a DOS prompt.

- On the client, execute the BIOS flash program.

In our example, we executed the program `DME521-010111.EXE`.

- When the BIOS flash program completes and before you reboot the client, on the frontend, reset the PXE action so when the client reboots, it will boot the OS from local disk:

```
# rocks set host runaction <hostname> action=os
# rocks set host boot <hostname> action=os
```

- When the flash process completes, reboot the client node.



Important Acknowledgement: We are grateful for the contribution of Jeremy Lawrence for describing the basic approach in his Rocks discussion list posting <https://lists.sdsc.edu/pipermail/npaci-rocks-discussion/2005-June/012092.html>.

8.4. Adding a Login Appliance to Your Cluster

A *login appliance* is a host that users use to build, launch and monitor their application(s). This host is a subset of a frontend appliance (a frontend runs several other services that are used to install and maintain all the nodes in the system).

First, we need to perform the initial installation of the login appliance. On the frontend, execute:

```
# insert-ethers
```

Select "Login" as the appliance type.

Then PXE boot the host that you want to be a login appliance.

Login hosts are generally hosts that also have a public network interface so users can directly access the login host. After the login appliance installs, to configure the public network interface on the login host, see Adding a public IP address to the second ethernet adapter on a compute node.

8.4.1. Making any Host a Submission and/or Execution Host

If you've installed a batch queuing system (e.g., by using the SGE Roll or Torque Roll), you can change any host in your system to be a submission host (a host where users can submit jobs) and/or an execution host (the queuing system schedules jobs on execution hosts).



By default, a login host is a submission host, but not an execution host. By default, a compute node is an execution host, but not a submission host.

To make a host a submission host (e.g., tile-0-0), on the frontend execute:

```
# rocks set host attr tile-0-0 submit_host true
```

Then reinstall tile-0-0.

To make a host an execution host (e.g., tile-0-0), on the frontend execute:

```
# rocks set host attr tile-0-0 exec_host true
```

Then reinstall tile-0-0.

8.5. Channel Bonding Interfaces

Channel bonding enables two or more network interfaces to act as one. The following procedure describes how to channel bond interfaces in Rocks.

We'll use the example of channel bonding two interfaces for compute-0-1. First, let's list the current configuration of the interfaces for compute-0-1:

```
# rocks list host interface compute-0-1
SUBNET  IFACE  MAC                IP                NETMASK          MODULE  NAME          VLAN  OPTIONS
private eth0   00:1e:4f:b0:74:ef  10.1.255.253     255.255.0.0      tg3             compute-0-1   ----  -----
----- eth1   00:10:18:31:74:43  -----          -----          tg3             -----      ----  -----
```

Here's what we'll do:

- Bond eth0 and eth1.
- Name the bonded channel "bond0".
- Give bond0 the IP address 10.1.255.253.
- Associate bond0 with the private network.

```
# rocks add host bonded compute-0-1 channel=bond0 interfaces=eth0,eth1 ip=10.1.255.253 network=pr
```

The above command changes the interface configuraion for compute-0-1 in the following way:

```
# rocks list host interface compute-0-1
SUBNET  IFACE  MAC                IP                NETMASK          MODULE  NAME          VLAN  OPTIONS  CHANNEL
private bond0  -----          10.1.255.253     255.255.0.0      bonding       compute-0-1   ----  -----  -----
----- eth0   00:1e:4f:b0:74:ef  -----          -----          tg3             -----      ----  -----  bond0
----- eth1   00:10:18:31:74:43  -----          -----          tg3             -----      ----  -----  bond0
```

To apply the configuration, execute:

```
# rocks sync config
# rocks sync host network compute-0-1
```

8.5.1. Adding Parameters to the Kernel Bonding Module

There are several options that can be added to the kernel bonding module (see RedHat's documentation [The Channel Bonding Module](#) ⁵ for a full description of all the options).

To set options "miimon=100" and "mode=balance-rr" for the kernel bonding module associated with bond0, execute:

```
# rocks set host interface options compute-0-1 bond0 options="miimon=100 mode=balance-rr"
```



The "options" value (e.g., "miimon=100 mode=balance-rr") must be a space-separated list.

And now we see:

SUBNET	IFACE	MAC	IP	NETMASK	MODULE	NAME	VLAN	OPTIONS
private	bond0	-----	10.1.255.253	255.255.0.0	bonding	compute-0-1	----	miimon=100 mode=
-----	eth0	00:1e:4f:b0:74:ef	-----	-----	tg3	-----	----	-----
-----	eth1	00:10:18:31:74:43	-----	-----	tg3	-----	----	-----

To apply the options to the bond0 interface, execute:

```
# rocks sync host network compute-0-1
```

8.6. Frontend Central Server

A Central Server is a Rocks Frontend node that can kickstart other frontends and provide rolls over the network, in a process called WAN kickstart. All Rocks frontends have the ability to act as central servers.

The standard Rocks distribution (located under /export/rocks/install) contains a distribution suitable for WAN kickstart. The only steps you must take is to open "www" and "https" access on your frontend for the purpose of RPM package transfer. See [Enable WWW access](#).

8.6.1. Adding Rolls to Serve from a Central

You may wish to serve rolls from your central server that you have not installed on your central server when the central server was initially installed. All frontends will serve the rolls they were built with to client frontends, but often it is advantageous to serve other rolls as well.

First, you must download the Roll ISO image to your central. Then, execute:

```
# rocks add roll <rollname>*.iso
```

8.7. Cross Kickstarting

Rocks supports heterogeneous clusters that contain nodes of different hardware architectures with a process called cross-kickstarting. To support an architecture different than its own, a frontend needs to expand its local distribution with additional packages. This section describes how to install distributions for other architectures on your frontend.

Start with a frontend node, as described by [Install Frontend Rocks 6](#), or [Upgrade Frontend](#). Follow the instructions below for every desired architecture.

8.7.1. Cross Kickstarting x86_64 Compute Nodes with a i386 Frontend

For this example, we assume the frontend is an i386 (32-bit) and the compute nodes are x86_64 CPUs (64-bit).

1. Download the x86_64 Rocks rolls. At a minimum, you'll need to download the x86_64 versions of the Base, Kernel, Web Server and OS disk 1 and OS disk 2 rolls.

For each roll, add it to the frontend's roll repository:

```
# rocks add roll <roll-name>.iso
# rocks enable roll <roll-name>
```

2. Rebuild your distribution:

```
# cd /export/rocks/install
# rocks create distro
# rocks create distro arch=x86_64
```

3. To prepare the frontend to cross-kickstart compute nodes via PXE, you first need to add the `rocks-boot` package from the x86_64 architecture:

```
# rpm -i --force --ignorearch /export/rocks/install/rocks-dist/x86_64/RedHat/RPMS/rocks-boot-7
# cp /boot/kickstart/default/vmlinuz-7.0-x86_64 /tftpboot/pxelinux
# cp /boot/kickstart/default/initrd.img-7.0-x86_64 /tftpboot/pxelinux
```

4. Then you need to add a new bootaction:

```
# rocks add bootaction action="install x86_64" kernel="vmlinuz-7.0-x86_64" \
  ramdisk="initrd.img-7.0-x86_64" \
  args="ks ramdisk_size=150000 lang= devfs=nomount pxe kssendmac selinux=0 noipv6"
```

5. The next step requires you to PXE boot the cross-kickstarted node twice: the first is to discover the node, and a second time to install it.

On the frontend, run `insert-ethers` to discover the compute node:

```
# insert-ethers
```

Then, PXE boot the compute node. After the compute node is discovered by `insert-ethers`, kill `insert-ethers` (by hitting the F9 key), then turn off the compute node.

Now instruct the node to use the "install x86_64" profile when it PXE boots (the command below assumes the compute node's name is `compute-0-0`):

```
# rocks set host installaction compute-0-0 action="install x86_64"
```

Now PXE boot the compute node again. The compute node will now install itself with the x86_64 distribution.

8.7.2. Cross Kickstarting i386 Compute Nodes with a x86_64 Frontend

For this example, we assume the frontend is a x86_64 (64-bit) and the compute nodes are i386 CPUs (32-bit).

1. Download the i386 Rocks rolls. At a minimum, you'll need to download the i386 versions of the Base, Kernel, Web Server and OS disk 1 and OS disk 2 rolls.

For each roll, add it to the frontend's roll repository:

```
# rocks add roll <roll-name>.iso
# rocks enable roll <roll-name>
```

2. Rebuild your distribution:

```
# cd /export/rocks/install
# rocks create distro
# rocks create distro arch=i386
```

3. To prepare the frontend to cross-kickstart compute nodes via PXE, you first need to add the `rocks-boot` package from the i386 architecture:

```
# rpm -i --force /export/rocks/install/rocks-dist/i386/RedHat/RPMS/rocks-boot-7.0-1.i386.rpm
# cp /boot/kickstart/default/vmlinuz-7.0-i386 /tftpboot/pxelinux
# cp /boot/kickstart/default/initrd.img-7.0-i386 /tftpboot/pxelinux
```

4. Then you need to add a new bootaction:

```
# rocks add bootaction action="install i386" kernel="vmlinuz-7.0-i386" \
  ramdisk="initrd.img-7.0-i386" \
  args="ks ramdisk_size=150000 lang= devfs=nomount pxe kssendmac selinux=0 noipv6"
```

5. The next step requires you to PXE boot the cross-kickstarted node twice: the first is to discover the node, and a second time to install it.

On the frontend, run `insert-ethers` to discover the compute node:

```
# insert-ethers
```

Then, PXE boot the compute node. After the compute node is discovered by `insert-ethers`, kill `insert-ethers` (by hitting the F9 key), then turn off the compute node.

Now instruct the node to use the "install i386" profile when it PXE boots (the command below assumes the compute node's name is `compute-0-0`):

```
# rocks set host installaction compute-0-0 action="install i386"
```

Now PXE boot the compute node again. The compute node will now install itself with the i386 distribution.

8.8. Adding Kernel Boot Parameters

There are two scenarios in which you need to add extra kernel boot parameters: 1) during installation, and 2) during normal boot. For both scenarios, we will use the example of adding the kernel boot parameter: `"ucsd=rocks"`.

8.8.1. Adding Kernel Boot Parameters to the Installation Kernel

The boot *action* of a node is controlled by the Rocks command line. For example, the following lists what the nodes will do on their next PXE boot:


```
# rocks list host boot
HOST      ACTION
vi-2:     -----
compute-0-0: os
compute-0-1: os
```

We see that `compute-0-0` will boot the "os", that is, it will boot off its local disk. Another common action is "install" which means the node will boot the installation kernel and install itself on the node's next PXE boot.

To list all the possible boot actions for a node, execute:

```
# rocks list bootaction
ACTION      KERNEL      RAMDISK      ARGS
install:    vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
install headless: vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
memtest:    kernel memtest  -----
os:         localboot 0      -----
pxeflash:   kernel memdisk bigraw pxeflash.img  keeppxe
rescue:     vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
```

To change a node's boot action to "install", execute:

```
# rocks set host boot compute-0-0 action="install"
```

Then we see that correct boot action has been applied:

```
# rocks list host boot
HOST      ACTION
vi-2:     -----
compute-0-0: install
compute-0-1: os
```

Now to change the kernel boot parameters for the installing kernel, we'll need to add an action. We'll create a new action called "install ucsd" and add our boot flag "ucsd=rocks" to the end of the kernel boot parameters:

```
# rocks add bootaction action="install ucsd" kernel="vmlinuz-5.2-i386" ramdisk="initrd.img-5.2-i386"
args="ks ramdisk_size=150000 lang= devfs=nomount pxe kssendmac selinux=0 noipv6 ucsd=rocks"
```

We now see that "install ucsd" is an available action:

```
# rocks list bootaction
ACTION      KERNEL      RAMDISK      ARGS
install:    vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
install headless: vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
install ucsd:  vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
memtest:    kernel memtest  -----
os:         localboot 0      -----
pxeflash:   kernel memdisk bigraw pxeflash.img  keeppxe
rescue:     vmlinuz-5.2-i386  initrd.img-5.2-i386  ks ramdisk_size=150000 lang= devfs=no
```

To set the install action for `compute-0-0` to "install ucsd", execute:

```
# rocks set host installaction compute-0-0 action="install ucsd"
```

Then we see that correct boot action has been applied:

```
# rocks list host
HOST      MEMBERSHIP CPUS RACK RANK RUNACTION INSTALLACTION
vi-2:     Frontend  1   0   0   os       install
```

```
compute-0-0: Compute 2 0 0 os install ucsd
compute-0-1: Compute 2 0 1 os install
```

The next time compute-0-0 is installed, it will use the "install ucsd" profile.

8.8.2. Adding Kernel Boot Parameters to the Running Kernel

To get the current parameters that will be appended to the default kernel boot parameters, execute:

```
# rocks report host bootflags
rocks-168: dom0_mem=1024M
compute-0-0: dom0_mem=1024M
```

Then, to add our boot flag of "ucsd=rocks", execute:

```
# rocks set host bootflags compute-0-0 flags="dom0_mem=1024M ucsd=rocks"
```

To check to make sure the flag has been added to the database, execute:

```
# rocks report host bootflags
rocks-168: dom0_mem=1024M
compute-0-0: dom0_mem=1024M ucsd=rocks
```

To apply the flags to the running kernel, you'll need to reinstall the compute nodes. For example, you'll need to reinstall compute-0-0 in order to apply "ucsd=rocks" to its kernel boot parameters. After compute-0-0 is reinstalled, you can check that the new kernel boot parameters have been applied by logging into compute-0-0 and executing:

```
# cat /proc/cmdline
ro root=LABEL=/ dom0_mem=1024M ucsd=rocks
```

8.9. Controlling Bootloader Parameters

The default parameters for the GRUB bootloader is the master boot record on the first BIOS-enabled drive. However, sometimes more control is needed to support alternative configurations or to work around either anaconda (the installer) issues or platform BIOS problems. For Rocks version 6, the default parameters are

```
--location=mbr --timeout=5
```

These are encoded into the base.xml configuration node and is part of the base roll.

It is possible to completely override the default with the *Bootloader* attribute. This affects the generation of the node's kickstart file. The most common change is the *--driveorder* option to tell anaconda precisely which device is the first in the BIOS boot order. In the following example, we set the driveorder to "sda" for node compute-0-0.

```
# rocks set host attr compute-0-0 Bootloader value="--location=mbr --timeout=5 --driveorder=sda"
```

One can check the value as follows

```
# rocks list host attr compute-0-0 | grep Bootloader
compute-0-0: Bootloader --location=mbr --timeout=5 --driveorder=sda H
```

We see that `compute-0-0` has the Bootloader attribute set. The "H" indicates that this attribute is a Host-specific attribute.



If you define a Bootloader attribute, you must include the `--location=` directive. This is an Anaconda requirement



The full set of bootloader options is available at Kickstart Options⁶

To revert to the Rocks default, one simply removes the attribute

```
# rocks remove host attr compute-0-0 attr=Bootloader
```

8.10. System Update

There are two ways in which packages can be updated on a Rocks Cluster. You can create an update roll and then add it to the distribution (preferred method) or you can update your frontend using YUM and then add the update packages to your distribution. The next two chapters will explain in detail the two methodologies.

8.10.1. System update using update roll

This approach is preferred over the other because it allows to test the effect of updates on compute nodes and in case of problems there is an easy way to roll back the system to the previous configuration.

1. Use `rocks create mirror` to create an update roll

```
# rocks create mirror http://mirror.centos.org/centos/6/updates/<arch>/Packages/ rollname=Centos_6_X_update_2014_03_21
```

Where *arch* is your current architecture (x86_64, i386) and the X should be the current minor release number (if latest stable version of Centos is 6.5 X should be 5).

2. If the current minor release number changed since you installed your Rocks Cluster or if it changed since last upgrade using this procedure you must also create a new OS roll relative to the current release minor. For example if the current release number of Centos is 6.5 run the following command:

```
# rocks create mirror http://mirror.centos.org/centos/6/os/<arch>/Packages/ rollname=Centos_6_X
```

If you are in doubt, you should run also the previous command, since the only drawback is that you will waste some disk space. X should be substituted with the current minor release.

3. Add the update roll created by the previous command to your distribution

```
# rocks add roll Centos_6_X_update_2014_03_21-6.1-0.x86_64.disk1.iso
# rocks add roll Centos_6_X-6.1-0.x86_64.disk1.iso
# rocks enable roll Centos_6_X
# rocks enable roll Centos_6_X_update_2014_03_21
# cd /export/rocks/install
# rocks create distro
```

4. New installed nodes will automatically get the updated packages. It is wise to test the update on a compute nodes to verify that updates did not break anything. To force a node to reinstall, run the command:

```
# ssh hostname-0-0 /boot/kickstart/cluster-kickstart
```

5. If something goes wrong you can always revert the updates removing the update roll.

```
# rocks remove roll Centos_6_X_update_2014_03_21
# rocks remove roll Centos_6_X
# cd /export/rocks/install
# rocks create distro
```

6. After you tested the update on some nodes with the previous step, you can update the frontend using the standard yum command

```
# yum update
```

Then reboot the frontend to make the update effective.

8.10.2. System update using YUM

To update using YUM you first have to update packages on your frontend (for this reason this approach is more risky) and then you can propagate the modification to the rest of the cluster.

1. Several repositories are configured, but then disabled by Rocks at install time. This is done so that the user can have full control over the distribution. The local distribution IS a yum repository. If you have created an updates roll as in the previous section, then

```
# yum clean all
# yum check-update
```

might produce output that resembles the following

```
[root@ ~]# yum check-update
Rocks-6.0 | 1.9 kB 00:00
Rocks-6.0/primary | 1.1 MB 00:00
Rocks-6.0 2183/2183
keyutils.x86_64 1.4-3.el6 Rocks-6.0
keyutils-libs.x86_64 1.4-3.el6 Rocks-6.0
keyutils-libs-devel.x86_64 1.4-3.el6 Rocks-6.0
[root@ ~]#
```

Alternatively, you can enable the CentOS updates repo and do the following

```
# yum clean all
# yum --enablerepo=updates check-update
```

might produce output that resembles the following

```
[root@ ~]# yum --enablerepo=updates check-update
updates | 3.5 kB 00:00

firefox.x86_64 10.0.4-1.el6.centos updates
keyutils.x86_64 1.4-3.el6 Rocks-6.0
keyutils-libs.x86_64 1.4-3.el6 Rocks-6.0
keyutils-libs-devel.x86_64 1.4-3.el6 Rocks-6.0
libpng.x86_64 2:1.2.49-1.el6_2 updates
libpng-devel.x86_64 2:1.2.49-1.el6_2 updates
libsmbclient.x86_64 3.5.10-116.el6_2 updates
```

```

openssl.x86_64                1.0.0-20.el6_2.4            updates
openssl-devel.x86_64         1.0.0-20.el6_2.4            updates
qemu-kvm.x86_64              2:0.12.1.2-2.209.el6_2.4    updates
samba-winbind-clients.x86_64 3.5.10-116.el6_2            updates
xorg-x11-server-Xephyr.x86_64 1.10.4-6.el6_2.1.0.1.centos updates
xorg-x11-server-Xorg.x86_64  1.10.4-6.el6_2.1.0.1.centos updates
xorg-x11-server-common.x86_64 1.10.4-6.el6_2.1.0.1.centos updates
xulrunner.x86_64            10.0.4-1.el6.centos         updates
[root@ ~]#

```

2. Now, Actually install the updates on your frontend

```
# yum --enablerepo=updates update
```

3. When YUM installs update packages on your system, it stores a copy of downloaded RPMs in the directory `/var/cache/yum/updates/packages`. It is necessary to add this packages to the distribution copying them in the contrib directory:

```
# cp /var/cache/yum/updates/packages/* /export/rocks/install/contrib/7.0/<arch>/RPMS/
```

Where *arch* is your current architecture (x86_64, i386).

4. Finally rebuild your distribution:

```
# cd /export/rocks/install
# rocks create distro
```

5. New installed nodes will automatically get the updated packages. To force a node to reinstall, run the command:

```
# ssh hostname-0-0 /boot/kickstart/cluster-kickstart
```

8.11. Reconfigure a Cluster

This section explains how to reconfigure a cluster after the initial installation, when some of the cluster wide configurations need to be changed (e.g. public IP, FQDN of the frontend, etc.). The reconfiguration process generates a script which performs all the necessary steps to modify the desired parameters. The script is generated traversing the graph and parsing the section marked as "reconfigure" section (see the developers guide).



Reconfiguring a cluster is a potentially harmful operation and can damage you system. The safest option is always reinstalling the frontend with the correct values.

8.11.1. Supported reconfiguration

Currently only the following roll support reconfiguration:

1. Base
2. Kernel
3. Ganglia

4. Java
5. Perl
6. Python
7. Hpc
8. SGE

The attribute which can be reconfigure are:

1. Kickstart_PublicHostname
2. Kickstart_PublicDNSDomain
3. Kickstart_PrivateHostname
4. Kickstart_PublicAddress
5. Kickstart_PublicBroadcast
6. Kickstart_PublicNetmaskCIDR
7. Kickstart_PublicNetwork
8. Kickstart_PublicNetmask
9. Kickstart_PublicGateway
10. Kickstart_PrivateAddress
11. Kickstart_PrivateBroadcast
12. Kickstart_PrivateNetmaskCIDR
13. Kickstart_PrivateNetwork
14. Kickstart_PrivateNetmask

The attribute names are self explanatory and they can be used to change public network configuration and private network configuration.

The reconfiguration works only if you haven't changed manually rocks default configurations (Postfix, Named, Apache, configuration files should still be at their default configurations)

8.11.2. How to reconfigure a system

To change one of the main parameters of your system first reset the history of your attributes change:

```
# rocks run reconfigure clear=1
```

To change the attributes, use the standard rocks commands (in the following example we show how to change the public ip address):

```
# rocks set attr Kickstart_PublicAddress 123.123.123.123
# rocks set attr Kickstart_PublicNetmask 255.255.255.0
# rocks set attr Kickstart_PublicGateway 123.123.123.1
# rocks set attr Kickstart_PublicHostname newhostname.somedomain.com
```

When you change an attribute (e.g. Kickstart_PublicAddress) the old value will be stored in an attribute named with a "_old" postfix (e.g. Kickstart_PublicAddress_old). You can verify which attributes will be updated with reconfigure using showattr:

```
# rocks run reconfigure showattr=1
User modified attributes
ATTR                                OLD VALUE ->                                NEW VALUE
Kickstart_PublicNetmask             255.255.255.0                                255.255.255.0
Kickstart_PublicHostname             myoldhost.ucsd.edu                            newhostname.somedomain.com
Kickstart_PublicAddress              2.2.2.24                                      123.123.123.123
Kickstart_PublicGateway              2.2.2.1                                       123.123.123.1

Additional attributes
Kickstart_PrivateHostname            myoldhost                                      newhostname
Kickstart_PublicBroadcast            2.2.2.255                                     123.123.123.255
Kickstart_PublicNetwork              2.2.2.0                                       123.123.123.0
Kickstart_PublicDNSDomain            ucsd.edu                                       somedomain.com
Kickstart_PublicNetmaskCIDR         24                                             24
```

The command shows also the attributes that will be updated to maintain consistency (section "Additional attributes") in the rocks database. If all the modified values make sense you can proceed and run the reconfigure script:

```
# rocks run reconfigure > /tmp/script.sh
# # inspect the script.sh
# bash /tmp/script.sh
```

Depending on what attributes you have changed you will need to reinstall all your nodes.

```
# rocks run host compute login "/boot/kickstart/cluster-kickstart-pxe"
```



If you have SGE installed the reconfigure operation will wipe all your SGE configurations (queue configuration, host configuration, etc.)



If you change some private networking parameters you will have to remove and re-install all your nodes with `insert-ether`. For example if you change the private network address (e.g. from 10.1.0.0 to 10.5.0.0) all the IP addresses of your node will be out of the new network. For this reason you will need to remove all the compute nodes and re-add them with `insert-ethernet`.

Notes

1. <http://code.google.com/p/google-authenticator>
2. <https://play.google.com/store/apps/details?id=com.google.android.apps.authenticator2>
3. <http://itunes.apple.com/us/app/google-authenticator/id388497605?mt=8>
4. http://en.wikipedia.org/wiki/QR_code
5. http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/Deployment_Guide/s2-modules-bonding.html
6. https://access.redhat.com/knowledge/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Installation_Guide/s1-kickstart2-options.html

Chapter 9. Command Reference

9.1. No Commands Installed at BUILD

install rocks-command-base

Appendix A. Frequently Asked Questions

A.1. Installation

1. Insert-ethers never sees new compute nodes. I also don't see any DHCP messages from compute nodes on the frontend. What is wrong?

Try bypassing the network switch connecting your nodes to the frontend. The switch may be configured to squash broadcast messages from unknown IP addresses, which drops DHCP messages from nodes. To verify your switch is indeed the problem:

1. Connect an ethernet cable between a single compute node and the frontend's "eth0" interface.
2. Install the compute node normally (Install Compute Nodes). You should see the DHCP messages from the node at the frontend.

2. While trying to install a compute node, and when I plug a monitor into the compute node, I see the error message 'Error opening kickstart file /tmp/ks.cfg. No such file or directory' or I see a screen on the compute node asking me to select a language. What went wrong?

A compute node kickstart requires the following services to be running on the frontend:

1. dhcpd
2. httpd
3. mysqld
4. autofs

To check if httpd and mysqld are running:

```
# ps auwx | grep httpd
# ps auwx | grep mysqld
```

If either one is not running, restart them with:

```
# /etc/rc.d/init.d/httpd restart
```

and/or

```
# /etc/rc.d/init.d/mysqld restart
```

The autofs service is called 'automount'. To check if it is running:

```
# ps auwx | grep automount
```

If it isn't, restart it:

```
# /etc/rc.d/init.d/autofs restart
```

Finally, to test if the Rocks installation infrastructure is working:

```
# rocks list host profile compute-0-0
```

This should return a kickstart file.

And to see if there are any errors associated with constructing kickstart files on the frontend:

```
# rocks list host profile compute-0-0 > /dev/null
```

3. I successfully installed all the Rolls, but during the last stage after the machine reboots, the system hangs with the error: *GRUB Loading Stage2....* What went wrong?

This is an intermittent problem we've seen in the lab as well. The installation is fine, except that the grub installation program, for an unknown reason, did not run correctly.

Here is a workaround:

- Put the Rocks Boot Roll CD in the frontend and boot the frontend.
- At the boot prompt, type:

```
frontend rescue
```

- A screen will appear, click the *Continue* button.
- When you see the shell prompt, execute:

```
# chroot /mnt/sysimage
```

- Run the grub installation program:

```
# /sbin/grub-install `awk -F= '/^#boot/ { print $2 }' /boot/grub/grub.conf`
```

This should output something similar to:

```
Installation finished. No error reported.
This is the contents of the device map /boot/grub/device.map.
Check if this is correct or not. If any of the lines is incorrect,
fix it and re-run the script `grub-install`.
```

```
# this device map was generated by anaconda
(fd0)      /dev/fd0
(hd0)      /dev/hda
```

- Exit the chroot environment:

```
# exit
```

- Reboot the frontend.
- Take the CD out of the drive and the frontend should come up cleanly.

4. When I try to install a compute node, the error message on the compute node says, "Can't mount /tmp. Please press OK to restart". What should I do?

Most likely, this situation arises due to the size of the disk drive on the compute node. The installation procedure for Rocks formats the disk on the compute node if Rocks has never been installed on the compute node before.

The fix requires changing the way Rocks partitions disk drives. See Partitioning for details.

5. My compute nodes don't have a CD drive and my network cards don't PXE boot, but my compute nodes do have a floppy drive. How can I install the compute nodes?

You will create a boot floppy that emulates the PXE protocol. This is accomplished by going to the web site:

ROM-o-matic.net¹

Then click on the version number under the *Latest Production Release* (as of this writing, this is version 5.4.3). Select your device driver in item 1. Keep the default setting in item 2 (Floppy bootable ROM Image). Then click "Get ROM" in item 4.

We suggest using `dd` to copy the downloaded floppy image to the floppy media. For example:

```
# dd if=eb-5.4.0-pcnet32.zdisk of=/dev/fd0
```

Then run `insert-ethers` on your frontend and boot your compute node with the floppy.

A.2. Configuration

1. How do I remove a compute node from the cluster?

On your frontend end, execute:

```
# rocks remove host "[your compute node name]"
```

For example, if the compute node's name is `compute-0-1`, you'd execute:

```
# rocks remove host compute-0-1
# rocks sync config
```

The compute node has been removed from the cluster.

2. Why doesn't startx work on the frontend machine?

Before you can run `startx` you need to configure XFree86 for your video card. This is done just like on standard Red Hat machines using the `system-config-display` program. If you do not know anything about your video card just select "4MB" of video RAM and 16 bit color 800x600. This video mode should work on any modern VGA card.

3. I can't install compute nodes and I have a Dell Powerconnect 5224 network switch, what can I do?

Here's how to configure your Dell Powerconnect 5224:

You need to set the *edge port* flag for all ports (in some Dell switches is labeled as *fast link*).

First, you'll need to set up an IP address on the switch:

- Plug in the serial cable that came with the switch.
- Connect to the switch over the serial cable.

The username/password is: admin/admin.

- Assign the switch an IP address:

```
# config
# interface vlan 1
# ip address 10.1.2.3 255.0.0.0
```

- Now you should be able to access the switch via the ethernet.
- Plug an ethernet cable into the switch and to your laptop.

- Configure the ip address on your laptop to be:

```
IP: 10.20.30.40
netmask: 255.0.0.0
```

- Point your web browser on your laptop to 10.1.2.3
- Username/password is: admin/admin.
- Set the *edge port* flag for all ports. This is found under the menu item: *System->Spanning Tree->Port Settings*.
- Save the configuration.

This is accomplished by going to *System->Switch->Configuration* and typing 'rocks.cfg' in the last field 'Copy Running Config to File'. In the field above it, you should see 'rocks.cfg' as the 'File Name' in the 'Start-Up Configuration File'.

4. How do I export a new directory from the frontend to all the compute nodes that is accessible under /home?

Execute this procedure:

- Add the directory you want to export to the file `/etc/exports`.

For example, if you want to export the directory `/export/disk1`, add the following to `/etc/exports`:

```
/export/disk1 10.0.0.0/255.0.0.0(rw)
```



This exports the directory only to nodes that are on the internal network (in the above example, the internal network is configured to be `10.0.0.0`)

- Restart NFS:

```
# /etc/rc.d/init.d/nfs restart
```

- Add an entry to `/etc/auto.home`.

For example, say you want `/export/disk1` on the frontend machine (named *frontend-0*) to be mounted as `/home/scratch` on each compute node.

Add the following entry to `/etc/auto.home`:

```
scratch frontend-0:/export/disk1
```

- Inform 411 of the change:

```
# make -C /var/411
```

Now when you login to any compute node and change your directory to `/home/scratch`, it will be automounted.

5. How do I disable the feature that reinstalls compute nodes after a hard reboot?

When compute nodes experience a *hard* reboot (e.g., when the compute node is reset by pushing the power button or after a power failure), they will reformat the root file system and reinstall their base operating environment.

To disable this feature:

- Login to the frontend

- Create a file that will override the default:

```
# cd /export/rocks/install
# cp rocks-dist/arch/build/nodes/auto-kickstart.xml \
site-profiles/7.0/nodes/replace-auto-kickstart.xml
```

Where *arch* is "i386" or "x86_64".

- Edit the file `site-profiles/7.0/nodes/replace-auto-kickstart.xml`
- Remove the line:

```
<package>rocks-boot-auto</package>
```

- Rebuild the distribution:

```
# cd /export/rocks/install
# rocks create distro
```

- Reinstall all your compute nodes



An alternative to reinstalling all your compute nodes is to login to each compute node and execute:

```
# /etc/rc.d/init.d/rocks-grub stop
# /sbin/chkconfig --del rocks-grub
```

A.3. System Administration

1. How do I use user accounts from an external NIS server on my cluster?

While there is no certain method to do this correctly, if necessary we recommend you use "ypcat" to periodically gather external NIS user accounts on the frontend, and let the default 411 system distribute the information inside the cluster.

The following cron script will collect NIS information from your external network onto the frontend. The login files created here will be automatically distributed to cluster nodes via 411. This code courtesy of Chris Dwan at the University of Minnesota.

```
(in /etc/cron.hourly/get-NIS on frontend)
```

```
#!/bin/sh
ypcat -k auto.master > /etc/auto.master
ypcat -k auto.home > /etc/auto.home
ypcat -k auto.net > /etc/auto.net
ypcat -k auto.web > /etc/auto.web

ypcat passwd > /etc/passwd.nis
cat /etc/passwd.local /etc/passwd.nis > /etc/passwd.combined
cp /etc/passwd.combined /etc/passwd

ypcat group > /etc/group.nis
cat /etc/group.local /etc/group.nis > /etc/group.combined
cp /etc/group.combined /etc/group
```



There is no way to insure that UIDs GIDs from NIS will not conflict with those already present in the cluster. You must always be careful that such collisions do not occur, as unpredictable and undefined behavior will result.

Notes

1. <http://www.rom-o-matic.net/>

Appendix B. Release Notes

B.1. Release 6.2 - changes from 6.1.1

B.1.1. Enhancements and Bug Fixes

Base: rocks run roll now supports host flag to specify the hostname

Base: rocks create distro now spawns 8 worker threads to improve speed of distro rebuild

Base: rocks add host verifies hostname is not an appliance name

Base: fix reverse path filtering on centos 6.x

Base: rocks set host bootflag now is applied to the local grub kickstart configuration

Base: security fix in kickstart.cgi and setPexboot.cgi

Base: rocks add appliance now supports specification of a distribution

Base: 411-alert-handler properly handles corrupted timestamp file

Base: creates rst files for roll commands

Base: properly apply MTU when using VLAN on an interface

Base: commands: fix report interface with IPMI interface

Base: do not propagate group with GID < 500 through 411

Base: set host ip now accepts NULL

Base: insert-ethers now verifies that required services are running at startup (dhcpd, httpd, tftpd, channeld)

Base: rocks create distro sets proper umask before creating the distro

Base: login and frontend nodes install fail2ban

Base: it is now possible to reconfigure networking of frontend without reinstallation (public/private ip setting, fqdn, gateway, etc.)

Base: <file> tag supports disabling RCS file generation

Base: <file> tag supports include attribute to simplify xml scripting

Base: firewall now supports SNAT

Base: Use NFSv4 for home directories and /share

KVM: CPU capabilities can be specified with attributes

KVM: support building virtual frontends on vm-containers when the container has a public interface

KVM: several plugins allow flexible configuration in several parts of KVM roll (starting a VM, stopping a VM, creating libvirt xml)

KVM: physical container can suspend and resume virtual machine when shutting down and when restarting

KVM: Virtual Machine can now mount ISOs as cdrom drive

KVM: virtual machine can now run on different appliances type (not just VM Container)

KVM: airboss client can now fetch private keys from ssh-agent

HPC: openmpi uses the proper BTL

ZFS: updated to 0.6.4.1 Support for the ZFS event Daemon is automatic.

ZFS: Bug Fix Properly rebuild ZFS modules when kernel is updated.

HTCondor: updated to 8.2.8

HTCondor: fixed UID_DOMAIN setting to be the local domain

HTCondor: configuration now supports Condor's config directory structure.

B.2. Release 5.6/6.1_sp1 - changes from 5.6/6.1

B.2.1. Enhancements and Bug Fixes

Base: rocks yum repository is clean from errors (automatic dependencies/provides detection has been disabled on several packages, in particular: all the opt-perl-* rpms, foundation-git, etc.)

Base: /opt/rocks/lib/mysql is not included in the loader path so system mysql uses system libmysqlclient.so and not rocks mysql libs (several other useless entries were removed from the ld.so.conf).

Base: JDK updated

Base: grub boot timeout is now 10 seconds

Base: improved script to fetch 441 private key after re-installation of a node

Base: it is now possible to use a /etc/ssh/ssh_known_hosts.local to insert static entries

Base: rocks set host interface ip does not allow anymore invalid IP number

Base, Ganglia, KVM: several fixes and improvement to the documentations

Base: Bugfix. record in rocks data base partitioning information

Base: rocks report host interface was failing with bonded interface

Ganglia: rrd uses previous data retention policy (as in ganglia 3.2) to reduce disk usage (ganglia 3.3 policy was to 39MB per node <http://marc.info/?l=npaci-rocks-discussion&m=135850438309730>).

B.3. Release 5.5/6.1 - changes from 5.5/6.0

B.3.1. New Features

- Host-based SSH authentication is now the default. This eliminates the requirement the users have password-less ssh keys and/or mounted home area on remote nodes
- Two-factor SSH authentication using Google Authenticator Apps for Android and/or iPhone is supported for any/all users. Users can easily create a time-based verification codes that are used in addition to their usual password.
- libqrencode.so and qrencode added to easily generate QR 2-dimensional codes.
- New ZFS-linux roll to support the ZFS file system via the ZFS on Linux¹. Does not support zfs as root file system, yet.

- New kernel-org roll to more easily support vanilla linux kernels from kernel.org.
- Added Infiniband Support (Mellanox Hardware). Subnet manager runs by default on frontend. Devices properly configured for users to run OpenMPI applications.

B.3.2. Enhancements and Bug Fixes

Base: removed creation of md5sum of all rpm packages during rocks create distro

Base: now properly configure rsyslog on Rocks 6 (instead of syslog)

Base: properly configure logrotate on Rocks 6

Base: better handling of "rocks run host" when running command on the local host

Base: added support for frontend with single partition

Base: new documentation section on ipmi and insert-ethers command line

Base: support for GPT partition table on the compute nodes (disk bigger than 2TB)

Base: Bugfix. Secure Attributes resynchronized after reinstall/reboot.

Base: Bugfix. Fixed "remove appliance" to raise an error if a host with that appliance type still exists.

Base: Bugfix. Support software RAID-1 for compute node / partitions

Base: Bugfix. For automatic reinstalls, grub configuration modified to add ksdevice= directive

Base: Bugfix. Fixed permissions on file: /etc/security/ca/ca.cfg

Base/Kernel: Bugfix. Kernel roll can now be rebuilt if just the Rocks-trimmed OS Roll is installed. Full CentOS distribution not required.

Kernel: tracker client now checks received rpm with internal rpm md5sum

Ganglia: updated Ganglia to 3.3.7

B.4. Release 5.5/6.0 - changes from 5.4.3

B.4.1. New Features

- Same Code base for CentOS 5 and CentOS 6 Systems

While there are significant differences between CentOS 5 and CentOS 6, Rocks now supports both from the same source code. 5 and 6 Rolls are *not interchangeable*.

- KVM Roll (virtualization) for CentOS 6 64-bit Systems

Xen is no longer supported by RedHat/CentOS 6 system and KVM (Kernel Virtual Machine) is only supported for 64-bit machines. KVM is hardware-only virtualization and requires both capability and BIOS to be enabled.

- Environment Modules

CentOS 6 natively supports environment modules instead of *-selector. For Rocks 5.5, the identical version of modules is compiled and included in the base roll.

- New support for Yum repositories

Standard OS-defined Repositories are no longer removed, but are instead disabled. the command `yum repolist all`, describes configured repos.

- Selectable Public and Private Networks at Install

Users are no longer required to select eth0 as the private and eth1 as the public networks for frontend installation. Install screens have drop-down menus to select the detected interface.

- Support of BIOS-named Ethernet devices

Version 6 kernels name Ethernet devices differently from Version 5. Devices may have names like p2p1, p35p2, and/or eth0. If you prefer standard device names for 6, you must install a node using the standard kernel flag of `biosdevname=0`

- Support of Frontends with One Physical Interface

The installer will configure a logical Ethernet device (e.g. eth0:0) if only one physical network device is detected during frontend installation.

- Creation of Upstart Scripts for CentOS 6

CentOS 6 use upstart to define to boot characteristics

- New Command: rocks report post

`rocks report post` will convert the post section of an xml file to an executable shell script. This is an advanced user feature.

- New Package: foundation-sqlite

This implements the SQLite3 single file database. See Upgraded Tracker.

B.4.2. Enhancements and Bug Fixes

Base: Updated anaconda

Base: Built against CentOS 5.8/6.2 with Updates as of May 8, 2012

Base: Java Updated to Version 7 Update 3

Base: Eclipse Moved to Java Roll and Updated to version 3.7.1

Base: Profile scripts only add once to PATH variables

Base: `rocks run host localhost` performs a fork instead of an ssh.

Base: Utilize subprocess and hashlib for newer version of Python

Base: Improve performance of report host dhcpd by nearly a factor of 10 on large clusters.

Base: Added Git RPM.

Base: Really fixed Zombie Processes in 411 service

Base: Patched foundation-python to remove a requirement on `/usr/local/bin/python`.

Base: Updated Version of M2Crypto to 0.21.1

Base: Better Bootstrapping support for building Rocks. Supports "bootstrap0" which starts on a bare OS install. Use yum whenever possible in package installation.

Base: Removed rebuild of Kudzu. CentOS fixed their internal build problems.

Base: Only have 32-bit or 64-bit version of Java. Not both. Works around broken Oracle-supplied RPMs where only one can be installed.

Base: Fixed various graph ordering issues.

HPC: Use environment modules. Set rocks-openmpi as default MPI Module

HPC: Added PVM RPMs for Version 6. PVM dropped from Native OS.

HPC: OpenMPI compiled with SGE integration

Ganglia: Updated profile.d scripts to only add once to PATH

Ganglia: rrdtool renamed to foundation-rrdtool to avoid conflicts with external repositories.

Ganglia: Removed a Package Conflict with ganglia-pylib

Ganglia: Insured that ganglia links with Rocks-compiled rrdtool

Java: Eclipse updated to 3.7.1

Java: JBOSS updated to 6.1.0. Version 5 was deprecated and would not build on CentOS 6

Kernel: Rewrote the rocks tracker to use SQLite for storing hashes. Eliminated segmentation faults of older code.

Kernel: handle difference between syslog and rsyslog

Kernel: build keyutils only on 5, not on 6.

SGE: Added Security Patched to Open Grid Scheduler v6.2 update 5p3

SGE: Fix compilation to work on both 5.8 and 6.2

Bio: Many build fixes for Perl modules.

Bio: Extracted profile scripts to an RPM for easier management

Bio: mpi-blast compiled with the rocks-opemmpi module loaded

Bio: BioPerl now depends on Perl Roll

Bio: BioPerl CPAN support utils updated

Condor: Update Condor to v7.6.6

Condor: No support for CentOS 6/32-bit.

Condor: Install simple condor test.

B.5. Release 5.4.3 - changes from 5.4

B.5.1. New Features

- Rocks Security

The root password supplied during the frontend installation is now used only for the root password of the frontend. Earlier, this password would be used for the MySQL database, Wordpress, and as root passwords for all compute nodes. All these services are now configured randomly-generated passwords.

To set the root passwords for individual backend nodes, the user can now use command line tool `rocks add host sec_attr ... attr=root_pw` to add a specific root password. `rocks sync host sec_attr ...` for the password to be set. Root passwd on the frontend is changed using the OS-supplied `passwd` command.

The rationale behind setting random root passwords for all backend nodes is that, if by some means, an attacker gained access to the root account of a backend node, and then the adversary could ran an offline attack against the crypted version of the root password, none of the other nodes would be compromised.

We've introduced the concept of secure attributes in the database. Any attribute such as passwords, private keys, etc can be stored in the secure attributes table in the database. This table is locked down to be readable only by root, and no one else. It's contents are not transferred during kickstart, and can only be synced using `rocks sync host sec_attr`, which will update the host specific secure attributes. Each attribute must also have a plugin associated with it, to specify the action to be performed on the backend nodes with the secure attribute as its data.

- Changes to 411 infrastructure

411 files can now only be requested by privileged accounts on backend nodes. This is enforced by checking that 411 requests originate only from privileged ports. Iptables is used to filter out requests that come in from non-privileged ports.

411 filters now support pre-send, filter, and post-receive functions. These functions are used to modify the contents of the 411 files being distributed, and to perform local system actions once the 411 files have been received and written to disk. These filters are present in the form of 411 plugins which are stored, modified, and enhanced on the frontend.

The 411 shared key is now distributed outside the kickstart file. This is driven by the command line. The client makes an RPC request to the frontend, and the frontend transfers the 411 shared key out-of-band. It uses rocks command line tools to verify that the request is coming from within the cluster.

- Introduction of the categories and indicies for resolving host specific properties in the database.

With this release, we've laid the foundation for creating and using random categories of host groups.

In the previous releases of rocks, the only categories available to system administrators by which they could group hosts were - Global category, Appliance category, OS category, and Hosts category.

With this release, we've introduced the capability of being able to create user specified categories. Some examples are - rack category, where hosts are grouped by rack, a bio category - where hosts can be grouped by whether or not a set of hosts have the bio roll installed, etc.

We've also introduced the concept of category resolution, where when resolving all the categories that a host can belong to, we can specify an chain of resolution. For example - we can state that compute-0-0 belongs to categories [global, linux, compute, rack0, bio]. In our resolution, we can state that we want the properties of the hosts to be picked up from [global, compute, rack0]. This way compute-0-0 picks up only the properties that are part of its resolution chain of categories.

Since this feature is still prototypical, at the moment, it is used only internally for firewall commands and single resolution chain. Subsequent releases will apply the same technique to attributes and routes.

- Changes to Firewall commands

The rocks firewall commands now require the presence of a rulename for every iptables rule. These rules are then ordered lexicographically by rule name.

The firewall command structure has some significant changes to reflect the categories and indicies feature described in the previous bullet point. Please look at the rocks command line documentation for more details on how to run the rocks firewall commands.

- Introduction of Perl and Python rolls

Two new rolls have been added to Rocks.

- Perl Roll - The Perl roll contains Perl 5.14.1, and plenty of CPAN modules that are required for the application software in Rocks to function properly. This version of Perl is installed in `/opt/perl/`.
- Python Roll - The version of Python that the core rocks utils depend on, is version 2.4.2. This is a rather dated version of Python. To provide users with the latest versions of python, we've create a Python roll which contains Python version 2.7.2 and version 3.2.1. These are both installed in `/opt/python/`.

B.5.2. Enhancements and Bug Fixes

Base: Updated anaconda to v11.1.2.224

Base: Built against CentOS 5.6 with Updates as of August 7, 2011

Base: database secured during installation. DB security now setup using script.

Base: Only root can create tables in cluster db. Use the rocks.my.cnf file to connect to db as root, because the db is already secured by the time this xml file.

Base: Default random root passwords for client nodes

Base: Added a Development Appliance. This is a backend appliance designed for building rolls without impacting the frontend. The appliance uses the frontend yum repository by default but can be configured to use its own local repository for full isolation.

Base: Removed foundation-perl, cpan, and cpan-support. These have been moved to the Perl roll

Base: Save the debug files from the ramdisk onto /root of the hard disk, so we can use them for post-install analysis.

Base: Properly report disk partitions into the database for software RAID file systems. Increase installation speeds for clients with software RAID file systems.

Base: 411 shared key is no longer transferred through kickstart. It is transferred through "rocks sync host sharedkey" 411 configuration is now generated through "rocks report host config411" 411 files are not transferred during kickstart. They are now transferred at first boot.

Base: Unambiguous add host command. Previously it was not possible add hostnames that were command line actions. For example, you could not add a host called "attr" because the command "rocks add host attr" exists.

Base: Remove shadow attributes from attributes tables and added secure attributes tables.

Base: Set primary interface of login servers back to private. Otherwise SGE behaves badly.

Base: Firewall rules now have rulenames and lexical ordering.

Base: Add profile.d/ssh-keygen to login appliances

Base: Hard link /etc/ssh/authorized_keys/id_rsa.pub to /root/.ssh/id_rsa.pub. If a user (or update) sets root's directory permissions tightly, we still read the public key.

Base: Block non-privileged traffic to 411 port from all networks, including localhost.

Base: Add variable to manage number parallel instances of "rocks sync host" commands.

Base: Support for pre-send, filtering, and post receive actions in 411.

Base: Minor modification to 411put. Use a get_filename function instead of a filename constant.

Base: 411 plugins can access host attributes

Base: 411 filters user password, and shadow information of users with UID < 500

Base: Initial Support for host categories and indicies.

Base: named.conf bug fixes - Now supports multiple subnets on non-octet boundaries.

Base: If yum install fails due to dependency error, force install using rpm --nodeps.

Base: Rocks run roll now honors the "--interpreter" flag to the post sections.

Base: Honor .<arch> directive to yum install. When installing packages use, "yum install <package>" instead of "yum install <packagefile>.rpm"

Base: Use YUM instead of RPM for rocks run roll This fixes two issues - On 64bit we were not installing the 32bit RPMs, and name.arch packages were not being installed.

Base: Support for HVM when using the Xen roll.

Base: Runtime optimization: Do not regenerate ALL pxeboot files. Just those for the hosts specified on the command line.

Base: Parallel class now takes care of serializing tasks that may overrun the system. If more than a set number of tasks are running, then requesting tasks will wait till slots are available to run Each task now prints out error messages if they fail on remote hosts. This way, we can track which syncs failed and which ones succeeded.

Base: Root ssh key needs to be passwordless to allow command/sync access to backend nodes. If root, create ssh key without passphrase. If normal user, create key interactively.

Base: Updated Java to 1.6 update 26

Base: Support for versioned centrals

HPC: Removed ganglia-web-frontend-addons from HPC roll

HPC: Update IOZone to 3.397

HPC: Update MPICH2 to v1.4

HPC: Update OpenMPI to v1.4.3

Ganglia: Updated to v3.2.0

Ganglia: gmond.conf cleaned to support updated version.

Ganglia: Updated apr to v1.4.5

Ganglia: Updated apr-utils to v1.3.13

Java: Tomcat-connectors rpm bug fix. Now no longer generates conflicts when installing tomcat httpd configuration.

SGE: Upgraded SGE to Open Grid Scheduler v6.2 update 5p2

SGE: Move the login appliance configuration out of the SGE roll and into the Base roll.

Web Server: Updated Wordpress to 3.1.3

Web Server: Updated Rocks Theme for Wordpress.

Web Server: Scrub root password from the installation and set the admin password to a string that never hashes.

Web Server: Wordpress admin password can be reset only if valid admin email is supplied.

Xen: Now support HVM as well as paravirtual instances

Xen: no longer need 'rocks-create-vlan'

Xen: added a report to create the xendomains configuration file

Xen: save the CA key and CA certificate that are used to authenticate

libvirt messages. Xen: touch /var/lock/subsys/xendomains in order to save running VMs

Xen: Ability to put frontend on arbitrary vm-container and set its name

Xen: Explicitly state default for virtualization type

Area51: set the right attr for the default tripwire email address

Area51: send tripwire reports to multiple recipients

Bio: Login appliance gets Bio Roll

Bio: Biopython now depends on Python Roll

Bio: BioPython upgraded to v1.5.7

Bio: BioPerl now depends on Perl Roll

Bio: BioPerl CPAN support utils updated
Bio: All BioPerl CPAN utils now built with cpan2dist
Bio: Updated EMBOSS to v6.3.1
Bio: Updated Autodock suite to v4.2.3
Bio: Update CGView to v2.0
Bio: Update fasta to v36.3.5a
Bio: Update MrBayes to v3.1.2
Bio: Update Blast to v2.2.25
Bio: Update reportlab to v2.5
Bio: Update t_coffee to v8.99
Bio: Update WGS to v6.1
Condor: Update Condor to v7.6.2
Condor: rocks login appliance submits jobs to condor
Condor: Experimental: Support submission to EC2
Condor: Add RANK parameter

B.6. Release 5.4 - changes from 5.3

B.6.1. New Features

- Redesign of the Avalanche Installer.

While observing the performance of the Avalanche Installer on a 1000-node machine, it became obvious that we must reduce as much traffic to the frontend as possible. This led to replacing the python BitTorrent-based installer with a BitTorrent-inspired installer written in C. The C code allows us to put more files into the peer-to-peer network, most notably: product.img (160KB), stage2.img (108MB) and updates.img (98MB).

To further reduce traffic to the frontend, the frontend now sends package predictions to installing nodes. When a node asks for a package, the tracker on the frontend sends a list of node addresses where that package can be found, plus a list of the next 9 packages that node will most likely ask for next. When similar appliances are concurrently installing, this reduces tracker traffic by 10x.

Installing nodes can be grouped. When an installing node asks the tracker for the location of a package and if other nodes are concurrently installing, the tracker will favor nodes that are in the same group as the requesting node, that is, the list the tracker sends back to the installing node will have nodes from the same group as the installing node at the top of the list. The default grouping is by rack, but it can be controlled by the "coop" attribute. For example, if you would like to put all nodes from rack 0 and rack 1 in the same group (named "red"), you would execute: "rocks set host attr rack0 rack1 coop red".

One can specify multiple trackers and multiple "package servers". A package server is a node that is "guaranteed" to have the requested file (e.g., the frontend).

For every downloaded file, an MD5 checksum verification is performed. This detects the case where a peer may have corrupted a file and prevents the corrupted file from spreading into the peer-to-peer network.

- Channel bonding for nodes is now controlled by the Rocks command line.

Channel bonding configuration for a node is stored in the database and can be added, removed or modified with the Rocks command line (e.g., "rocks add host bonded ..."). After channel bonding is configured for a node, it can be dynamically applied by executing "rocks sync host network ...".

- All nodes' firewall rules are controlled by the Rocks command line.

The rules for all the nodes are stored in the database and can be added, removed or modified with the Rocks command line (e.g., "rocks open host firewall", "rocks close host firewall", "rocks remove host firewall"). After a node's firewall settings are changed, they can be applied to the node on-the-fly with "rocks sync host firewall 'hostname'" (this command is also called when the user executes "rocks sync host network ...").

- Introduction of "Air Traffic Control".

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 powering up, powering down and installing a virtual cluster can be controlled by the (authorized) non-root user.

In addition to VM power control, we've also added the ability to attach to a VM's console. This allows users to see the entire boot sequence for a VM starting from the "BIOS" boot messages.

Several Rocks commands were added to support this feature: "rocks create keys" (to create public/private key pairs), "rocks set host power" (to power up/down VMs and to forcibly install a VM, akin to PXE booting a physical machine), and "rocks open host console" (to attach to a VM's console).

- "greceptor" replaced with "channeld".

The wire protocol for Ganglia messages changed which required a major overhaul to greceptor. We made the decision to write a simple RPC-based service (named 'channeld') to take over the responsibilities of greceptor. Channeld accepts 411-put requests and acts on them by using 411-get to download files under the control of 411.

All other components of 411 remain unchanged, only the notification engine has been enhanced.

- DNS resolution for multiple domains.

The DNS naming system on the frontend now supports multiple zones, where each subnet managed by the frontend can be put into a different zone. The DNS service can be turned on or off for each individual zone.

- Login appliance support.

A node can be configured as a Login appliance. By default, a Login appliance can submit jobs, but it cannot execute jobs.

- Set the name of a host based on the name of a specific network interface.

The "primary_net" attribute allows nodes to have /bin/hostname set to the name of a network interface other than "private". This is useful for login or other multiple interface appliances.

- Easily swap 2 interfaces with one Rocks command.

To swap the settings of 2 interfaces, execute "rocks swap host interface ...".

- Created a GIT repository for Rocks-related source code.

The host "git.rocksclusters.org" is a GIT repository for all core Rocks code, UCSD Triton Resource code and Rocks contrib code.

B.6.2. Enhancements

OS: Based on CentOS release 5/update 5 and all updates as of November 2, 2010.

Base: Anaconda installer updated to v11.1.2.209.

Base : no longer remap the private network to "eth0", instead Rocks keeps track of the network a node kickstarted from and maps that network to the "private" network. For example, if a node kickstarted off "eth1", then "eth1" will be mapped to the private network.

Base : hardened the Anaconda installer to more aggressively write the grub configuration files onto the boot disk. This helps to mitigate the "hang while trying to load Grub stage2" issue.

Base : removed ext4 kernel module from installation environment. We found that trying to mount a swap partition as an ext4 file system frequently caused kernel panics during installations.

Base : added ksdevice=bootif to all the PXE boot targets. This improves installation speed by reusing the IP address/interface information when a node PXE boots. Previously, a node would re-scan all ethernet interfaces.

Base : when a node XML file has a syntax error, "rocks list host profile" prints out the name of the node XML file and the line number where the syntax error occurred.

Base : "rocks run host" now spawns multiple parallel threads when multiple hosts are supplied. Also added the following parameters: timeout (thanks Tim Carlson!), delay, stats, collate and num-threads.

Base : yum configuration default modified to bind to the frontend's public IP instead of the private. This facilitates easy package installation for external nodes (e.g., nodes running on a public cloud).

Base : non-existent attributes are considered to be false conditionals when building configuration files.

Base : "precedes" method added for Rocks command plugins to enable fine-grained ordering of plugin execution.

Base : network interfaces under Linux support 2 new specific modes: "dhcp" and "noreport". The "dhcp" mode indicates that the interface should always DHCP to get its address. The "noreport" mode specifies that no "ifcfg-*" file should be written for the interface. If a mode is not specified for an interface, then Rocks will create an "ifcfg-*" file for the interface based on values set in the database (just like it did in the previous release).

Base : IPMI now uses the interface channel column in the networks table to specify the baseboard controller channel number.

Base : text inside "changelog" tags is now wrapped in CDATA to allow XML escape characters. This is only supported for node XML files found within Rolls (not for node XML files found under /export/rocks/install/site-profiles).

Base : rolls can be built without a complete copy of the Rocks source code. They use the Rocks development environment found under /opt/rocks/share/devel on a frontend.

Area51: tripwire updated to v2.4.2.

Bio: refreshed CPAN modules.

Bio: refreshed CPAN MPI-Blast.

Bio: added Celera Whole Genome Sequence Assembler.

Condor: updated to v7.4.4.

Condor: automated Condor configuration completely retooled: 1) the configuration is Rocks command based instead of standalone CondorConf tool, 2) it supports dynamic update of any/all configurations on nodes, 3) it uses Rocks command plugins to allow additional automated condor config (e.g., via plugin, it can turn on MPI support).

Condor: supports a pool password (shared secret) for additional host verification.

Condor: integrates with EC2 roll to extend Condor pools with EC2 Hosts.

Condor: support added for port ranges to facilitate firewall configuration.

Condor: local copy of Condor's manpages added to roll documents.

Condor: support for updating Condor on nodes without re-installation (e.g., rocks run host "yum update condor" ; rocks sync host condor).

Ganglia: monitor-core updated to v3.1.7.

Ganglia: rrdtool updated to v1.4.4.

Ganglia: the Ganglia Roll can now be added on-the-fly to an existing frontend.

Ganglia: all nodes send out their metric metadata every 3 minutes. In the past, when gmond was restarted on the frontend, it couldn't collect metrics from the nodes because it had no metadata from the nodes (and it didn't have a way to ask the nodes because the nodes are configured in "deaf" mode).

HPC: iofzone updated to v3.347.

HPC: iperf updated to v2.0.5.

HPC: MPICH2 updated to v1.2.1p1.

HPC: OpenMPI updated to v1.4.3.

HPC: rocks-openmpi is the default MPI and it is configured with mpi-selector.

SGE: SGE updated to V62u5.

SGE: any host can be configured to be an execution host by setting the host's "exec_host" and "sge" attributes to true and any host can become a submission host by setting the host's "submit_host" and "sge" attributes to true.

Web-server: mediawiki updated to v1.16.0.

Web-server: wordpress updated to v3.0.1.

Xen: any node can now host Xen virtual machines. This is controlled with the "xen" attribute.

Xen: set the power for all nodes in a virtual cluster (except the VM frontend) with one command ("rocks set cluster power ..."). Power settings can be "on", "off" or "install" (turn on and force installation).

Xen: allow virtual machines to define VLAN tagged interfaces. Previously, VLAN tagging was only supported for physical interfaces.

B.6.3. Bug Fixes

Base: non-root users can no longer see the encrypted passwords with 'rocks list host attr'. Hashed passwords are now stored in a 'shadow' column in the attribute tables.

Base: the "%" in "rocks run host %" now returns all hosts. Thanks to Tom Rockwell for the fix.

Base: If an ethernet switch sends out a DHCP request, the DHCP server no longer sends it the "filename" and "next server" in the DHCP response. This caused some switches not to properly load their firmware. More generally, this is controlled by the "kickstartable", "dhcp_filename" and "dhcp_nextserver" attributes.

Base: "rocks set password" asks the user to confirm their new password.

Base: when a node requests a kickstart file and if the frontend determines that the frontend is too "busy", the kickstarting node now correctly does a random backoff before re-requesting its kickstart file. Prior to this fix, a node would backoff for 30 seconds.

Base: multiple conditionals can now be present in XML tags.

Base: fixed a graph traversal issue. In the past, if you had the graph "a" (cond) to "b" to "c" and if "cond" was false, the graph traversal would include "a" and "c". Now it just includes "a".

Base: permissions set in the "file" tag are preserved even if there are other "file" tags for the same file that don't set the file's permissions. The bug was when a later "file" tag without a "perms" attribute was encountered, the file's permissions were cleared.

Base: "file" tags now support "os" conditionals.

Base: in insert-ethers, appliances that are marked "not kickstartable" will not have to wait for a kickstart file. In the past, one had to hit the "F9" (force quit) key to exit insert-ethers when discovering non kickstartable appliances (e.g., ethernet switches).

Base: IPMI configuration cleaned up. Rocks no longer generates erroneous entries in modprobe.conf or /etc/sysconfig/ifcfg-ipmi.

Base: The "pre" tag now supports the "interpreter=" attribute.

Bio: eliminated "Permission Denied" errors during multiple runs on the same BLAST database by different users.

SGE: made the job collection metric more efficient. Previously, when 100's of jobs are submitted to a frontend's queue, the SGE metric would take so long to execute, it caused gmond to stop gathering metrics for all hosts.

SGE: the number of CPUs array jobs consume are now correctly counted.

B.7. Release 5.3 - changes from 5.2

B.7.1. Enhancements

OS: Based on CentOS release 5/update 4 and all updates as of December 15, 2009.

Area51: chkrootkit updated to v0.49.

Base: Anaconda installer updated to v11.1.2.195.

Base: Added support HP's Smart Array controllers (e.g., the 'cciss' driver) in the Rocks partitioning code.

Base: Moved JDK to the Base Roll (the Java Roll has been discontinued).

Base: Added the 'ssh_use_dns' attribute (default setting is 'true'). If set to 'false', then DNS is not consulted when trying to make an SSH connection. This speeds up SSH connections for frontend's that don't have access to a functioning DNS server.

Base: Added the 'Kickstart_DefaultLeaseTime' and 'Kickstart_MaxLeaseTime' attributes to give the user the ability to control their own DHCP lease timeouts.

Base: Added IMPI tools and configuration utilities.

Base: Updated foundation-perl to v5.10.1.

Base: Updated foundation-perl-Tk to v804.028.

Base: Experimental CPAN support. This introduces a utility that lets the user create Rocks RPMs directly from CPAN along with all dependencies.

Bio: Updated Emboss to v6.1.0.

Bio: Updated biopython to v1.52.

Bio: Updated clustalw to v2.0.12.

Bio: Updated fasta to v35.4.9.

Bio: Updated fftw to v3.2.2.

Bio: Updated gromacs to v4.0.5.
Bio: Updated iolib to v1.12.1.
Bio: Updated phylip to v3.69.
Bio: Updated t_coffee to v8.14.
Ganglia: Updated rrdtool to v1.3.8.
HPC: Updated iozone to v3.291.
HPC: Updated iperf to v2.0.4.
HPC: Updated MPICH2 to v1.1.1p1.
HPC: Updated OpenMPI to v1.3.3.
SGE: Updated to v6.2 update 4.
Viz: Updated nVidia driver to v190.42.
Web-server: Updated mediawiki to v1.15.1.
Web-server: Updated wordpress to v2.8.4.

B.7.2. Bug Fixes

Area51: tripwire is now correctly configured on the frontend.

Base: Increase the amount of memory PHP can use from 64 MB to 256 MB. This is essential for viewing Ganglia data for large clusters.

Base: During installation, don't copy over the rocks-cdrom block device file. On most systems, there is no issue, but there are some systems where this code would copy over the entire contents of the CD/DVD if the CD/DVD was still in the tray (which is shouldn't be, because after the rolls are copied to the frontend, the CD/DVD is unmounted and ejected).

Base: Set the '--utc' flag when hwclock is called. Without the flag, the localtime is written to the hardware clock. Then, at first boot, the system reads the hardware clock respecting the UTC settings in /etc/sysconfig/clock. So the system time ends up being off by the offset to UTC.

Base: On first boot, Set the number of CPUs for the frontend in the database to actual number of detected CPUs. In the past, this value was hard-coded to 1.

Base: Fix for frontend's that use DNS servers that are configured to return a 'catch all' address for non-existent domain names. Do a reverse lookup to get the IP address. In the previous release, one would see an 'XXX' entry for the IP address of the frontend in /etc/hosts.

Base: Put double quotes around 'option domain-name', otherwise, if the domain name was structured like xxx.yyy, then the DHCP service would not start.

Base: For tentakel, throttle the number of concurrent connections to 100 (useful for large clusters).

Base: Rewrite the pxelinux.cfg files after setting the run/install action in the nodes table (e.g., 'rocks set host installation ...').

Base: For 'rocks set password', make sure password changing code accesses the Rocks foundation database.

Base: For 'rocks sync host network', make sure to update the static routes file.

Base: Fixed 'insert-ethers --replace'. Before, it correctly removed the node, but then readded it at 'the end' of namespace, that is, it forgot the rack/rank of the replaced node.

Base: Make tentakel more resilient to hanging nodes.

Ganglia: Fix to save all RRD data on frontend shutdown.

Ganglia: Fixes to physical views. Make sure all nodes are positioned in rack/rank order.

Ganglia: Fixes to rediscover all backend nodes after a frontend reboot. Without this fix, if gmond was restarted on the frontend, then it requires restarting gmond on all the compute nodes in order to get metrics to report again.

Kernel: Fixes to preserve the timestamp from RPMs on Rolls that are installed via CD/DVD. Prior to this fix, all RPMs in Rolls that were installed from CD/DVD would have the timestamp of the time when the frontend was installed. This causes issues when trying to update RPMs after the frontend was installed.

SGE: When a host is removed (e.g., 'rocks remove host ...'), make sure the queue associated with that node is also removed.

SGE: Build the DRMAA library with Java support.

Xen: Respect the MTU value of the physical interface when bringing up Xen bridges.

Xen: Fix to enable physical devices (partitions, LVM partitions, etc.) to be used as devices for virtual disks.

Xen: Fix to 'ip' address flag in 'rocks add host vm'. Without the fix, if an IP address was specified on the command line, the command would throw an exception.

Xen: Fix to support for virtual compute nodes that are managed by a physical frontend.

B.8. Release 5.2 - changes from 5.1

B.8.1. New Features

- Solaris support for client nodes

With the new JumpStart Roll, one can now install and configure a Linux-based Rocks frontend to "JumpStart" Solaris-based back-end machines.

- Attributes

Can assign "attributes" to nodes at four levels: global, appliance type, OS (e.g., Linux or SunOS), and host. An attribute can be accessed in an XML node as an entity. For example, if you assign the attribute "foo" with the value "123" to compute-0-0 (i.e., with the command, "rocks set host attr foo 123"), then in an XML node file, you can access the value of the attribute foo with "&foo;"

Attributes also enable "conditionals". Using the example above, a "post" section can be optionally executed based on the value of an attribute. For example, if a post section is defined as: <post cond="foo='123'"> then this post section will only be executed if the attribute "foo" is set to "123" for the installing host.

B.8.2. Enhancements

OS: Based on CentOS release 5/update 3 and all updates as of June 22, 2009.

Base: Anaconda installer updated to v11.1.2.168.

Base: Isolated MySQL for the Rocks database under /opt/rocks.

Base: Converted all 'dbreports' to the Rocks command line.

Base: Configure 'MTU' for networks through the Rocks command line.

Base: Configure network routes through the Rocks command line.

Base: Configure host aliases through the Rocks command line.

Base: Added '/var/log/authpriv' to log rotate list.

Base: Added 'iburst' flag to NTP configuration on frontend.

Base: Added version and release info to vmlinuz and initrd.img to enable cross-kickstarting via PXE and making it easier to host different Rocks VMs on a physical system.

Base: Added a YUM configuration file to each node that points to the distribution on the frontend.

Bio: Updated biopython to v1.50.

Bio: Updated clustalw to v2.0.11.

Bio: Updated fasta to v35.4.7.

Bio: Updated fftw to v3.2.1.

Bio: Updated elph archive in glimmer to v1.0.1.

Bio: Updated gromacs to v4.0.1.

Bio: Updated perl-Data-Stag to v0.11.

Bio: Updated perl-Digest-MD5 to v2.38.

Bio: Updated perl-File-Temp to v0.21.

Bio: Updated perl-GD to v2.41.

Bio: Updated perl-GD-SVG to v0.33.

Bio: Updated perl-Graph to v0.91.

Bio: Updated perl-HTML-Parser to v3.60.

Bio: Updated perl-HTML-Tagset to v3.20.

Bio: Updated perl-PathTools to v3.30.

Bio: Updated perl-SOAP-Lite to v0.710.08.

Bio: Updated perl-SVG to v2.49.

Bio: Updated perl-SVG-Graph to v0.02.

Bio: Updated perl-Scalar-List-Utils to v1.21.

Bio: Updated perl-Storable to v2.20.

Bio: Updated perl-Text-Iconv to v1.7.

Bio: Updated perl-URI to v1.38.

Bio: Updated perl-XML-Parser to v2.36.

Bio: Updated perl-XML-Twig to v3.32.

Bio: Updated perl-XML-Writer to v0.606.

Bio: Updated perl-bioperl to v1.6.0.

Bio: Updated perl-libnet to v1.22.

Bio: Updated perl-libwww-perl to v5.826.

Bio: Updated phylib to v3.68.

Bio: Updated reportlab to v2.3.

Bio: Updated t_coffee to v7.81.

Ganglia: Updated to v3.1.2.

HPC: Updated OpenMPI to v1.3.2.

HPC: Updated MPICH2 to v1.0.8p1.

HPC: Updated stream to v5.9.

Java: Updated java to v1.6.0_13.

Java: Updated jdk to v6 update 13.

Java: Updated antlr to v3.1.

Java: Updated jboss to v5.0.1.GA.

Java: Updated jogl to v1.1.1.

SGE: Updated to v6.2 update 1.

Viz: Support for single, dual and quad display nodes.

Viz: Chromium support for 32-bit and 64-bit applications.

Viz: Added CUDA driver. Users can optionally run CUDA programs on the tiled-display nodes.

Viz: User updatable nVidia driver. Makes it easy for users to refresh the nVidia driver without having to wait for an updated Viz Roll.

Xen: Using libvirt instead of 'xm' command line programs to start/stop VMs.

Xen: Allow VM disk to be backed by a physical disk partition.

Xen: Use threading in the 'rocks add cluster' command to decrease the time to add a virtual cluster.

B.8.3. Bug Fixes

Base: Fix for software RAID partitioning.

Base: Increase timeout for package downloads when on slow networks (e.g., 100 or 10 Mbit).

Xen: Fix to ensure all routes are active after Xen is started.

Xen: Fix 'rocks set host vm' command to allow users to resize a VM's disk.

B.9. Release 5.1 - changes from 5.0

B.9.1. New Features

- Support for Virtual Clusters

Virtual frontends and virtual compute nodes are now supported. The network for a VM frontend its VM compute nodes are contained within its own VLAN.

A virtual cluster is added with "rocks add cluster fqdn=X ip=Y num-computes=Z". See "rocks add cluster help" for details.

- Can build rolls outside of Rocks source tree.

All roll building support files are under `/opt/rocks/share/devel`.

- Can reconfigure a compute node's network without rebooting.
Rocks commands were added to support this. See the documentation for the procedure.
- Distribution moved to `/export/rocks/install`.
No longer require NFS on the frontend to properly host a Rocks Distribution. This will make moving user accounts to an external NFS server easier.
- Fine-grained control over the "boot" and "install" kernel for Xen VMs.
Rocks commands were added to support this feature. For details, execute: "rocks help installprofile" and "rocks help bootprofile".

B.9.2. Enhancements

OS: Based on CentOS release 5/update 2 and all updates as of November 4, 2008.

Base: Anaconda installer updated to v11.1.2.113.

Base: Increased the / partition default size to 16 GB.

Base: Opened the 'www' and 'https' ports to the local public network.

Base: In Avalanche Installer, added code to check if a package is requested twice in a row. If it is, we assume the package is corrupted. In this case, we toss the package and retrieve it the package from the frontend.

Base: Added Rocks commands to manage the "aliases" table.

Base: Added "rocks remove roll" command. Thanks to Brandon Davidson from the University of Oregon for the code.

Base: The command "rocks-dist" is replaced with "rocks create distro".

Base: Disabled the watchdog for frontend installs that boot off a CD/DVD.

Base: Changed boot command from "frontend" to "build". To build a frontend, when you see the "boot:" prompt, now type: "build".

Web Server: Wordpress updated to v2.6.1.

Web Server: Updated Wordpress theme.

Area51: All commands converted to Rocks command line.

HPC: Updated OpenMPI to v1.2.7.

HPC: Updated MPICH2 to v1.0.7.

Java: Fixed a bug in the graph that Java from properly installing on compute nodes.

Restore: All files under `/export/rocks/install/contrib` are now included in the Restore Roll.

Restore: All files in `/var/named/*local` are now included in the Restore Roll.

Restore: The frontend's ssh machine keys are now included in the Restore Roll.

SGE: Updated to v6.1 update 5.

SGE: Added a script to reinstall a cluster by submitting an SGE job.

B.9.3. Bug Fixes

Base: Reverse domain lookups now work for subnets that don't fall on an octet boundary.

Base: Bootflags now carry over between reinstallations and reboots. In the previous release, a reboot would "forget" the bootflags set by the user with the Rocks command line.

Base: Added full path to "mksquashfs". Now can build distribution when using "sudo".

Restore: Ethernet Switches, Power Units and Remote Management appliances are now properly saved in the Restore Roll.

Ganglia: Fixed the 'tail +4' bug in the cron job.

SGE: Fixed the display in the "Job Queue" on the frontend's web site. SGE now reports the correct number of CPUs in use.

B.10. Release 4.3 - changes from 4.2.1

B.10.1. New Features

- Rocks Command Line

Initial release of the Rocks command line which facilitates non-SQL administrative access to the database. All Rocks commands have a regular structure of "rocks <verb> <component>". For example, to list all hosts that have been discovered by the frontend, execute: "rocks list host".

All rocks commands can be listed by executing: rocks. Also, help is included with each command. For example, for help on the command "rocks add host", execute: "rocks add host help".

For an overview of the Rocks command line, see Introduction to the Rocks Command Line². The reference for all Rocks commands can be found here³.

- PXE First

Hosts can now be configured in BIOS with a boot order of CD, PXE, Hard Disk (previous releases of Rocks required: CD, Hard Disk, PXE). In combination with the Rocks command line, node-specific installation parameters are easily supported. For details on PXE First, see Boot Order and PXE First⁴.

Note: The boot order of (CD, HD, PXE) continues to be supported in Rocks 4.3. That is, existing Rocks clusters can be upgraded without requiring the cluster owner to change any BIOS settings.

B.10.2. Enhancements

OS: Based on CentOS release 4/update 5 and all updates as of July 4, 2007.

Base: Anaconda installer updated to v10.1.1.63.

Base: Performance improvement when building torrent files for the Avalanche Installer.

Base: Database indirects. More flexibility with Rocks variables.

Grid: Globus updated to gt4.0.4 with web services.

Condor: updated to v6.8.5.

PVFS2: updated to v2.6.3.

Java: updated to v1.5.0_10.

Ganglia: updated to v3.0.4.

HPC: Now using OpenMPI and PVM from RedHat distribution.

B.10.3. Bug Fixes

Base: Install now supports machines which have more than 26 disk drives.

Base: 411 clients now atomically update files.

Condor: Max heap size properly set for java programs on small and large memory machines.

Condor: All logging written to /var/opt/condor.

B.11. Release 3.2.0 - changes from 3.1.0

New Feature - Added the Condor Roll. This brings the distributed high-throughput features from the Condor project to Rocks clusters.

New Feature - Added the Area51 Roll. This roll contains security tools and services to check the integrity of the files and operating system on your cluster.

New Feature - Ganglia RSS news event service.

Enhancement - Improved network handling for compute nodes: any interface may be used for the cluster private network, not simply the default "eth0".

Enhancement - Better support for cross-architecture clusters containing x86 and x86_64 machines.

Enhancement - GM device driver now builds and loads on compute nodes that have a custom kernel (e.g., a kernel from kernel.org).

Enhancement - Software RAID for custom compute node partitioning is supported.

Enhancement - Added variables for root and swap partition. If you only want to change the size of root and/or swap, you only have to reassign two XML variables.

Enhancement - The default root partition size has been increased to 6 GB (up from 4 GB).

Enhancement - SGE ganglia monitor added. The state of all SGE jobs can be tracked from the frontend's web page.

Enhancement - PXE support extended to support floppy-based Etherboot and ia64.

Enhancement - EKV uses ssh instead of telnet for security.

Enhancement - New Myrinet MPICH version 1.2.5..12.

Enhancement, Java Roll -- Updated JDK to version 1.4.2_04

Enhancement - Latest software updates recompiled for three architectures from RHEL source rpms.

Enhancement - Automatic MySQL Cluster database backup.

Enhancement - MAC addresses are included for each node in the "Cluster Labels" output.

Enhancement - Frontend rescue mode on the Rocks Base CD enabled. By typing "frontend rescue" at the boot prompt will give you a shell in which you can examine the state of the frontend.

Bug Fix - 411 hardened. More reliable notification of changed files. Correct Makefile encrypts login files on frontend first-boot.

Bug Fix - Multiple CD drives are supported for bringing up a frontend. If you have more than one CD drive connected to your frontend, the installer will now correctly identify which CD you are using.

Bug Fix - Ganglia metrics are now saved on frontend reboot. After a reboot, all Ganglia history will be restored from the previous boot.

Bug Fix - PVFS compiled with `-mcmmodel=kernel` on Opteron.

Bug Fix - XML escape characters (e.g., `&`, `<`, `>`) are supported in the installation screens (e.g., the Cluster Information screen and the Root Password screen).

Bug Fix, Intel Roll - All the Intel compiler libraries are now copied to the compute nodes.

B.12. Release 3.2.0 - changes from 3.1.0

New Feature - Added the Condor Roll. This brings the distributed high-throughput features from the Condor project to Rocks clusters.

New Feature - Added the Area51 Roll. This roll contains security tools and services to check the integrity of the files and operating system on your cluster.

New Feature - Ganglia RSS news event service.

Enhancement - Improved network handling for compute nodes: any interface may be used for the cluster private network, not simply the default "eth0".

Enhancement - Better support for cross-architecture clusters containing x86 and x86_64 machines.

Enhancement - GM device driver now builds and loads on compute nodes that have a custom kernel (e.g., a kernel from kernel.org).

Enhancement - Software RAID for custom compute node partitioning is supported.

Enhancement - Added variables for root and swap partition. If you only want to change the size of root and/or swap, you only have to reassign two XML variables.

Enhancement - The default root partition size has been increased to 6 GB (up from 4 GB).

Enhancement - SGE ganglia monitor added. The state of all SGE jobs can be tracked from the frontend's web page.

Enhancement - PXE support extended to support floppy-based Etherboot and ia64.

Enhancement - EKV uses ssh instead of telnet for security.

Enhancement - New Myrinet MPICH version 1.2.5..12.

Enhancement, Java Roll -- Updated JDK to version 1.4.2_04

Enhancement - Latest software updates recompiled for three architectures from RHEL source rpms.

Enhancement - Automatic MySQL Cluster database backup.

Enhancement - MAC addresses are included for each node in the "Cluster Labels" output.

Enhancement - Frontend rescue mode on the Rocks Base CD enabled. By typing "frontend rescue" at the boot prompt will give you a shell in which you can examine the state of the frontend.

Bug Fix - 411 hardened. More reliable notification of changed files. Correct Makefile encrypts login files on frontend first-boot.

Bug Fix - Multiple CD drives are supported for bringing up a frontend. If you have more than one CD drive connected to your frontend, the installer will now correctly identify which CD you are using.

Bug Fix - Ganglia metrics are now saved on frontend reboot. After a reboot, all Ganglia history will be restored from the previous boot.

Bug Fix - PVFS compiled with `-mcmmodel=kernel` on Opteron.

Bug Fix - XML escape characters (e.g., &, <, >) are supported in the installation screens (e.g., the Cluster Information screen and the Root Password screen).

Bug Fix, Intel Roll - All the Intel compiler libraries are now copied to the compute nodes.

B.13. Release 3.1.0 - changes from 3.0.0

Base Linux packages compiled from publicly available RedHat Enterprise Linux 3 Source (Advanced Workstation) for all architectures.

Switched to Sun Grid Engine 5.3 as the default batch scheduling system.

More Rolls: NMI/Globus Release 4, Java, Condor, Intel compiler rolls available.

New Architectures: Opteron (x86_64) receives first-class functionality.

Enhancement - New MPICH version 1.2.5.2. More efficient MPD parallel job-launcher handling. MPICH2 included by default as well.

Enhancement - Using latest Myrinet mpich-gm 2.0.8 for all architectures.

Enhancement - Updated SSH version 3.7.1 with no login delay.

Enhancement - 411 Secure Information Service used by default, replacing NIS.

Enhancement - Greceptor replaces Gschedule to support mpring, 411, cluster-top and others. Achieves an order of magnitude better performance than its predecessor.

B.14. Release 3.0.0 - changes from 2.3.2

Based on RedHat 7.3 for x86 and RedHat Advanced Workstation 2.1 for ia64 (all packages recompiled from publicly available source).

Enhancement - Includes RedHat updated RPMS (and recompiled SRPMs for ia64), as of September 3 2003.

Enhancement - Includes kernel version 2.4.20-20.7 for x86 and version 2.4.18e.37 for ia64. Installation environment includes all drivers from the above kernel packages.

Enhancement - New full-featured DNS server and structured ".local" naming conventions within cluster.

Enhancement - Linpack (xhpl) works out of the box for Pentium IV and Athlon.

Enhancement - Added remove node feature to `insert-ethers`.

Enhancement - New layout of all MPICH transports. See `/opt/mpich` on the frontend for the new directory structure.

Enhancement - Add support for 'Rolls'. An x86 Rocks frontend install now requires two CDs: the Rocks Base CD and the HPC Roll. An ia64 frontend still requires only one DVD.

Enhancement - Added 'Grid' Roll. This roll includes all packages from NMI R3.1, which includes Globus, the Simple Certificate Authority, and other packages.

Enhancement - High-Performance, fault-tolerant MPD job launcher made available. Automatic MPD ring creation and healing via KAgreement-mpd protocol. (Currently in beta phase for this release)

Enhancement - New 411 Secure Information Service to replace NIS. (Currently in beta phase for this release)

Enhancement - Latest Ganglia version 2.5.4 including better webfrontend speed and streamlined appearance, and more efficient network and disk metric handling.

Enhancement - New PhpSysInfo page on compute nodes, available along with /proc link on Ganglia host view page.

Enhancement - Ganglia command line tool has new --clustersize and --alive=host options.

Enhancement - Kickstart graph now viewable from frontend web page.

Enhancement - For kickstart graph files, new <file> tags made available, with owner="root.root" and perms="ga+r" attributes. Beta phase of RCS-based tracking of all config file changes made for post-section repeatability.

Enhancement - Kickstart graph ordering is explicit. Previously the evaluation order of individual nodes depended on graph weights. Node dependencies can now be explicitly specified using <order> tags in the graph files.

Bug Fix - UNIX manual pages correctly shown (we extend /etc/man.conf)

Bug Fix - NTP now synchronizes all compute node clocks with the frontend.

Bug Fix - add-extra-nic now supports multiple NICs per compute node.

Bug Fix - Ganglia RRD metric histories are archived on physical disk and restored on startup.

Bug Fix - Includes NCSA's OpenPBS scalability patches. Can now launch PBS jobs that require more than 64 processors.

Bug Fix - USB keyboard works on all ia64 Tiger boxes

B.15. Release 2.3.2 - changes from 2.3.1

Bug fix - Memory leaks in the broadcastSSH gmetric python module are fixed.

Bug fix - Gmetad will not crash when long ganglia metric names are introduced in the cluster.

Bug Fix - Building MPICH-GM package correctly for AMD Athlon processors.

Bug Fix - Added PBS directories: /opt/OpenPBS/sched_priv, /opt/OpenPBS/sched_logs, /opt/OpenPBS/undelivered.

Bug Fix - Added userdel that correctly updates the NIS database.

Enhancement - The Rocks-specific Ganglia metrics are much more efficient with a new Python C extension module that publishes ganglia metrics. The PBS job-queue monitor particularly benefits from this new module.

Enhancement - Updated rocks-boot package to contain all the modules from the latest kernel-BOOT package.

Enhancement - The Ganglia monitor-core and webfrontend packages have been updated to the latest version 2.5.3.

Enhancement - The frontend is now a fully configured Rocks cluster build host. By checking out all the Rocks source code on a 2.3.2 frontend, one can build all the source code simply by executing `make rpm` in the directory `.../rocks/src/`.

Enhancement - Updated SGE packages from v5.3p2-4 to v5.3p3-1.

Enhancement - Added Rocks version number to /home/install/contrib directory structure.

B.16. Release 2.3.1 - changes from 2.3

Bug fix - Now all the installation device drivers from Red Hat's device disks are included (e.g., Broadcom's Ethernet adapters). In Rocks 2.3, only the device drivers found on Red Hat's installation boot floppy were included.

Bug fix - User-specified NIS domains are now supported (in Rocks 2.3, only 'rocks' NIS domain was supported).

Bug fix - User-specified compute node disk partitioning is now supported.

Bug fix - Sun Grid Engine commd port errors during post installation and Sun Grid Engine warnings during `insert-ethers` were fixed.

Bug fix - Building for Pentium II/III and Athlon added to ATLAS RPM. (on a side note, ATLAS is now built against gcc version 3.2).

Enhancement - PVFS upgraded to version 1.5.6.

Enhancement - More detail has been added to the PBS queue monitoring web page (e.g., can view jobs for only one user and can view nodes for one job). Additionally, the monitoring code now more efficient and it has been hardened due to direct experiences on a 300-node Rocks cluster.

Enhancement - The `bssh` service has been moved from a standalone service to a task managed by the Ganglia `gschedule` service.

Enhancement - The ethernet-based MPICH package has been updated to version 1.2.5.

Enhancement - The Myrinet-based MPICH package has been updated to version 1.2.5..9.

Enhancement - OpenPBS version 2.3.16 has replaced PBS. Additionally, the *big memory* patch has been applied. Also, the license for OpenPBS requires registration for those that use OpenPBS, so if you use OpenPBS to manage your computational resources, please register at <http://www.OpenPBS.org>.

Enhancement - The `maui` package has been updated to version 3.2.5.

Enhancement - Updated Myricom's GM to version 1.6.3.

New Feature - Added a link of the main web page of the frontend that allows one to make sheets of labels with the names of all the compute nodes.

New Feature - An alternative version of `gcc` is now installed (version 3.2 is installed in `/opt/gcc32/...`).

B.17. Release 2.2.1 - changes from 2.2

Bug fix - `pvfs` and `gm` modules don't build because the kernel source and kernel binary RPMs were of a different version.

Bug fix - the partitioning on compute nodes only partitioned the first drive. Now all drives on compute nodes are partitioned with a single partition. The default partitioning is: 4 GB root partition, then `/state/partition1` is the remainder of the first drive. The second drive, if present, will have one partition labeled `/state/partition2`. The third drive, if present, will have one partition labeled `/state/partition3`, etc.

Bug fix - the Rocks CD didn't support as many hardware devices as the RedHat CD. All the hardware modules found on the RedHat CD have been added to the Rocks CD (including many, many more).

B.18. Release 2.2 - changes from 2.1.2

Based on RedHat 7.2.

Upgraded Ganglia (provided by Matt Massie of UC Berkeley) to 2.1.1.

Incorporated PVFS RPMs that were graciously provided to us from Najib Ninaba and Laurence Liew who work at Scalable Systems Pte Ltd in Singapore.

`insert-ethers` looks to see if a Rocks distribution exists. If it doesn't, `insert-ethers` rebuilds it.

Upgraded MPICH-GM to version 1.2.1..7b.

Added the "stream" memory bandwidth benchmark.

Added functionality to rocks-dist so distributions can be rebuilt without having to mirror the entire distribution.

Implemented a "greedy" partitioning scheme on compute nodes. The default partitioning is: 4 GB root partition, then /state/partition1 is the remainder of the first drive. The second drive, if present, will have one partition labeled "/state/partition2". The third drive, if present, will have one partition labeled "/state/partition3", etc.

Bug fix - added a "watchdog" timer to kickstart. This reboots a kickstarting node if it can't find a kickstart file. This problem was reported by folks trying to kickstart multiple nodes at the same time.

Bug fix - increased the polling intervals for maui so it won't time out when asking PBS about node status on larger clusters.

Bug fix - makedhcp now adds the full pathname to pxelinux.0 when it builds dhcpd.conf.

Bug fix - create a device node for /dev/cdrom.

Bug fix - /var/log/messages is now appropriately rotated.

B.19. Release 2.1.2 - changes from 2.1.1

Many network and storage drivers have been added to the installation CD. For example, SMC 83c170 EPIC/100 (epic100.o), RTL8139 SMC EZ Card Fast Ethernet (8139too.o) and the Promise SuperTrak Driver (pti_st.o) have all been included (as well as about 100 more).

The cluster configuration web form has been simplified.

The initial kickstart file that is generated from the web form is now streamed directly back to the user (rather than displaying the kickstart file, and then asking the user to save the file). This should finally kill the "I saved my kickstart file on Windows" problem.

An option to manually partition a frontend disk has been added to the cluster configuration web form.

The recursive directory /home/install/install/install/... has been eliminated.

Ganglia's axon is now started before pbs-server, as the pbs-server initialization script asks ganglia for the number of processor in each node when it creates one of it's configuration files.

The latest "stable" release of Myricom's GM (1.5) and MPICH-GM (1.2.1..7) packages.

High-Performance Linpack is now precompiled for Myrinet and Ethernet.

B.20. Release 2.1.1 - changes from 2.1

The main change in this release is the use of an XML-based *kickstart graph* to actively manage kickstart files.

Includes support for IA-64 compute nodes. See the Installing IA-64 Compute Nodes HOWTO⁵ for detailed information.

A full X server is now installed on frontend machines.

Added PXE support for kickstarting compute nodes.

All compute nodes now install ATLAS and high-performance Linpack -- some slick software from the Innovative Computing Laboratory⁶ at the University of Tennessee.

Modified to the PBS server initialization script to dynamically determine the number of CPUs in compute nodes by querying `ganglia`.

Created a `rocks-pylib` package that contains all the common code used by Rocks command line utilities that access the MySQL database, thus giving all the tools the same basic functionality and common user-specified flags.

Patched Red Hat's installation tool (`anaconda`) so the default behavior is to get kickstart files with HTTP (Red Hat's default is NFS). This frees the installation procedure of requiring NFS for *any* of its functions.

Rewrite of `insert-ethers` to give it the look and feel of a standard Red Hat installation tool.

Now using Red Hat's `pump` instead of `dhclient` for the DHCP client.

Properly create the default PBS configuration file (`/usr/apps/pbs/pbs.default`) so PBS is now operational "out of the box".

Fixed the annoying, but harmless, message `"socket.error: (101, 'Network is unreachable')"` that was seen on frontend boots.

Fixed the annoying, but harmless, message `"user 0 unknown"` that was seen on a compute node's first boot after kickstarting.

Fixed the 444 permissions problem on `/usr/man` and moved all the Rocks man pages into the new home for Linux man pages (`/usr/share/man`).

B.21. Release 2.1 - changes from 2.0.1

The main change in this release is that thanks to RedHat 7.1, we now use the Linux 2.4 kernel.

Based on RedHat 7.1, instead of 7.0.

Linux 2.4.x kernel, instead of 2.2.x.

Cluster-dist has been replaced with Rocks-dist. Command line arguments are very similar, with the `explode` command being removed and replaced with the `--copy` flag. The new Rocks-dist creates smaller distributions, fixes the problem of expensive mirror updating, and simplifies CD building. Also, it no longer deletes the distribution before rebuilding, this means the build directory (where kickstart files reside) is persistent across distribution builds.

Frontend is now a stratum 10 NTP server, so compute nodes will clock sync to the frontend even when the frontend cannot reach an external time source.

Usher daemon now correctly daemonizes, since we patch the GM code to allow processes to fork.

Symbolic links for Ekv and piece-pipe RPMs removed from the build directory, and `"@Control@"` section added to kickstart files.

`Pbs_mom_config.h` generated in the kickstart build directory.

Added pre-defined types to the models table in the SQL database. Also, removed dead tables from database, and made column order more human friendly.

Add SQL parsing to cluster-[ps]killfork scripts.

Removed cluster-config-compute, and cluster-config-frontend from the `"%post"` section in the kickstart file. The cluster-config rpm is now build and installed on the fly on each compute-node.

Bumped lilo timeout to 5 seconds.

Added `FORCE_UNIPROCESSOR` macro test to force sick SMP machines to kickstart as uniprocessor nodes.

Major revision of `insert-ethers`. Can now be used to replace nodes, and start at arbitrary ranks and basenames.

Minor maui and pbs bug fixes.

Added `gm-mpich SHMEM` support to `mpi-launch`.

B.22. Release 2.0.1 - changes from 2.0

Changed to new directory structure according to RedHat. Existing users will have to delete their mirror of www.rocksclusters.org and re-mirror to pickup the current RedHat directory naming scheme. NOTE: you need the new cluster-dist from www.rocksclusters.org to create a new mirror!

Added support to kickstart laptops (still working on this)

Frontend can now have either a DHCP or static address for the external network. For DHCP the DNS information provided from the external DHCP server is inserted into the Rocks Database and propagated to compute nodes.

Increased default DHCP lease time

Replaced Linux's useradd with create-account.

Force glibc-common RPM to be installed. RedHat 7.0 doesn't install this due to errors in the RPM database.

NIS database gets rebuilt on the frontend once an hour.

Create directories on frontend/compute nodes before putting down SSL and SSH keys. Fixed permission on directories.

Ssh-agent now forwards through nodes

Ssh doesn't use privileged port (makes firewalls happy)

cluster-kickstart set real and effect UID to root so all members of the install group can run shoot-node. Previously only root could do this.

Fixed reinstalls on IDE and SCSI hosts (only IDA host worked before, thanks to a RedHat 7.0 change)

Fixed bssh bug

Notes

1. <http://zfsonlinux.org>
2. <http://www.rocksclusters.org/roll-documentation/base/4.3/commandline.html>
3. <http://www.rocksclusters.org/roll-documentation/base/4.3/c229.html>
4. <http://www.rocksclusters.org/roll-documentation/base/4.3/boot-order.html>
5. ../howto/ia64.php
6. <http://icl.cs.utk.edu/>

Appendix C. 411 Secure Information Service Internals

The 411 Secure Information Service provides NIS-like functionality for Rocks clusters. It is named after the common "411" code for information in the phone system. We use 411 to securely distribute password files, user and group configuration files and the like.

411 uses Public Key Cryptography to verify file contents, and shared key cryptography to protect transport. It operates on a file level, rather than the RPC-based per-line maps of NIS. 411 does not rely on RPC, and instead distributes the files themselves using HTTP (web service). Its central task is to securely maintain critical login/password files on the worker nodes of a cluster. It does this by implementing a file-based distributed database with weak consistency semantics. The design goals of 411 include scalability, security, low-latency when changes occur, and resilience to failures.

C.1. Using the 411 Service

The 411 system intentionally mimics the NIS interface for system administrators. Of course there are elements in 411 which are not present in NIS, namely RSA cryptographic keys. However we have attempted to make 411 as easy to use as an NIS replacement.

Files listed in `/var/411/Files.mk` are automatically serviced by 411. This means that any file listed there will be kept up to date by the 411 agents on all compute nodes in your cluster. This is done using the makefile `/var/411/Makefile` in a similar fashion to NIS. To force the 411 system to flush all changes, execute the following on the frontend node:

```
# make -C /var/411
```

Note that this command is run by cron every hour on the frontend to propagate changes to compute nodes. New files can be added to `Files.mk` as necessary for custom services on the cluster.

To force all 411 files to be re-encrypted and change alerts sent to all compute nodes, run this on the frontend

```
# make -C /var/411 force
```



The 411 service uses IP broadcast messages on your cluster's private network to achieve optimal performance.

To force all compute nodes to retrieve the latest files from the frontend, execute:

```
# rocks run host command="411get --all"
```

C.2. Structure

C.2.1. Listener

The 411 system uses a Rocks RPC service called `channel` to communicate with client nodes. The client nodes run the RPC service, and the frontend has a RPC client which sends out alerts on its multicast channel. When the clients receive a 411-alert RPC message, along with the name of the file that changed, it does a pull of the file using `411get`.



411 is akin to a distributed database, and is not a centralized lookup service like NIS. While scalable, 411 does not provide instantaneous distribution of new files. The delay between running the 411 `makefile` and all nodes receiving the changed file depends on cluster size. A large password file on a cluster with many nodes can take up to a minute to fully synchronize on all nodes.

C.3. 411 Groups

Beginning in Rocks 3.3.0, 411 has the ability to send messages to subsets of the cluster. This facility, called 411 groups, allows us to distribute different files to nodes depending on their type. The group mechanism depends on the client nodes specifying group names in their local 411 configuration file; these are called the client's "registered" groups.



There is no per-group key in 411. The groups mechanism is only a convenience feature, without strong security to enforce it. Specifically, a node can eavesdrop on messages for a foreign group that it is not a member of.

Group names are multi-level, and resemble file paths. By default, every node is a member of the `'/'` group (corresponding to the traditional top-level 411 group), and the `'/Membership'` group, where *membership* is the node membership in the frontend database with the spaces replaced by underscores, so compute nodes are part of "Compute" and NAS Appliance are part of the "NAS_Appliance" group (this is due to a limitation of `makefile` in handling spaces).



A special Makefile called `/var/411/Group.mk` is available to help you setup and maintain 411 groups. After editing this file to specify which files go to which group, run

```
# make -C /var/411 groups
# make -C /var/411
```

To activate the 411 group `makefile` actions.

By default, nodes are members of a group with the same name as their *Membership*. For example compute nodes are automatically a member of the group "Compute". A sample `411.conf` file with several groups looks like:

```
<!-- Configuration file for the 411 Information Service -->
<config>
<master url="http://10.1.1.1:372/411.d/" score="0"/>
<group>/light/blue</group>
<group>Compute</group>
</config>
```

Multi-element group names have a simple inheritance model: specific groups imply more general ones. For example, if you are a member of the group `/compute/light`, you will automatically be interested in messages in group `/compute/light` and `/compute`. You will not be interested in messages from group `/compute/heavy`. In this case `/compute/light` is the specific group, and `/compute` is the more general one.

C.4. Plugins

Starting with Rocks 5.4, file specific plugins have been introduced into the 411 system. These plugins manipulate the content of files before sending files on the frontend, and after reception of the file on the client nodes. The plugins also mangle name of the file, ownership and mode of the file.

C.4.1. Plugin API

The 411 Plugin architecture follows a very simple API. Each plugin is written in Python and the plugins reside in `/opt/rocks/var/plugins/411/`. All plugins inherit the `rocks.service411.Plugin` class.

1. `get_filename`: This function returns a the filename on the frontend on which this plugin will function. This is the only required function in the plugin. All other functions are optional.

```
import os
import sys
import stat
import rocks.service411

class Plugin(rocks.service411.Plugin):
    def get_filename(self):
        return '/etc/auto.master'
```

2. `filter_name`: This function mangles the filename of the file at the destination, and returns the mangled name

```
def filter_name(self, fname):
    if fname == '/etc/auto.master' && self.attrs['os'] == 'sunos':
        return '/etc/auto_master'
```

3. `filter_owner`: This function takes the Owner UID and group ID of the file as a string argument and returns a mangled version of the "UID.GID" string. This string is in the format that the **chown** command understands.

```
def filter_owner(self, oid):
    if self.attrs['os'] == 'linux':
        return oid
    if self.attrs['os'] == 'sunos':
        return '0.0'
```

4. `filter_mode`: This function takes the mode information as a string, and returns a mangled mode information string. This should be in the numerical format that the **chmod** command understands.

```
def filter_mode(self, mode):
    return '010644'
```

5. `filter_content`: This function takes the contents of the file as a string, manipulates it and returns the final string to be stored in the file. The example below illustrates insertion of a blank line between every line of the input content.

Example C-1. `filter_content`

```
def filter_content(self, content):
    new_content = []
    for line in content.split('\n'):
        new_content.append(line + '\n' + '\n')
    return ".join(new_content).strip()
```

6. `pre_send`: This function manipulates the contents of the file before the file is made available for download over 411. The example below illustrates deletion of all blank lines from the content.

Example C-2. `pre_send`

```
def pre_send(self, content):
    new_content = []
    for line in content.split('\n'):
        if line.strip() is "":
            continue
        else:
            new_content.append(line)
    return '\n'.join(new_content).strip()
```

7. `post`: This function is run after the file is received, filtered and written to disk. The example below illustrates restarting the autofs service after a file has been written to disk.

Example C-3. `post`

```
def post(self):
    os.system("service autofs restart")
```

C.5. 411get Configuration File

As mentioned above, 411get configuration is stored in a file on the client nodes. This file is generated by a rocks command.

```
#rocks report host config411 compute-0-0
<file name="/etc/411.conf" perms="0600" owner="root:root">
<![CDATA[<!-- 411 Configuration -->
<config>
```

```
<appliance>compute</appliance>
<master url="http://10.1.1.1:372/411.d/" />
<os>linux</os>
<group>Compute</group>
</config>]]>
</file>
```

The configuration file contains information about the client node, stored as XML entities. For eg. - the appliance, group, and OS information. This entities are converted to attributes that are made available to the plugins. When 411get is run, the config file is read, and the OS, appliance, etc. information is stored in a dictionary called `self.attr`. This dictionary is made available to all the plugins, so that the plugins may filter based on the attribute of the nodes.

The admin may extend the config file to include entities other than the default ones. For this a rocks command line plugin must be created that outputs XML tags. Please see how to create a rocks command line plugin in the Rocks Developers' Guide.

Once the rocks command line plugin is created for the **rocks report host config411** command, the configuration can be synced to the client nodes using the following commands.

```
# rocks report host config411 compute-0-0 | ssh compute-0-0 "rocks report script | sh"
```

C.6. Commands

C.6.1. 411get

```
411get [--all] [--master=url] [--conf] [--pub] [--shared] [--local] [file]
```

Retrieves and decrypts 411 messages. Prints resulting file to stdout. When invoked with no files, 411get will list the available 411 messages.

The following options are available:

- *--all* Retrieves and writes all available 411 messages from the most attractive master. Does not print output to stdout, nor ask for confirmation before overwriting files.
- *--master* The url of a 411 master server to use. Defaults to "http://10.1.1.1:372/411.d/" or whatever is present in "/etc/411.conf". If given, this master takes precedence over those listed in the configuration file.
- *--file*, *--local* Assume the file is local, ie present in the current directory. Does not use http to retrieve the file. Decrypts and prints the file contents.
- *--conf* The configuration file to use. Defaults to "/etc/411.conf".
- *--pub* The location of the cluster public RSA key. Defaults to /etc/411-security/master.pub.
- *--shared* The location of the cluster shared key. Defaults to /etc/411-security/shared.key

The master servers, along with their quality score, are listed in the /etc/411.conf file on compute nodes.

C.6.2. 411put

```
411put [--411dir=dir] [--urldir=dir] [--see] [--noalert] [--alert=channel]
[--411name] [--pub] [--priv] [--comment=char] [--chroot=dir]
[--chroot-here] [--group=group] file1 file2 ...
```

Encrypts and publishes files using the 411 secure information service. Will send a broadcast message to client nodes by default, alerting them of a changed file.

The following options are available:

- *--chroot=dir* Turn "dir" into the root directory of the destination file. This allows files to be located in a different place on the master and clients.

Example:

```
411put --chroot=/var/411/groups/compute /var/411/groups/compute/etc/passwd
```

Will put "/var/411/groups/compute/etc/passwd" on compute nodes as "/etc/passwd".

- *--chroot-here* A convenience option, equivalent to *--chroot=\$PWD*.
- *--group=name* A 411 group for this file. Clients will ignore 411 messages in groups which they are not a part of. Allows 411 files to be published to a subset of the cluster. Name is path-like: "Compute/green", or "/Compute/green". Spaces are ok: "a space/yellow" is a valid group name as well.
- *--comment* The comment character for this file. Used to place a descriptive header without disrupting normal operations. Often set to "#". Default is none.
- *--411dir* The local directory to place encrypted 411 messages. Defaults to "/etc/411.d/". Be careful about the permissions of this directory.
- *--urldir* The web directory where 411 messages are available. Defaults to "/411.d/".
- *--see* Shows the encrypted file contents on stdout.
- *--noalert* Suppresses alert message.
- *--411name* Prints the 411 message name for the file. Provided for convenience.
- *--pub* The location of the 411 master public RSA key. Defaults to a 1024 bit key in "/etc/411-security/master.pub". This file should have permissions 0444 (read by all) and be owned by root.
- *--priv* The location of the 411 master private RSA key. Defaults to a 1024 bit key in "/etc/411-security/master.key". This file should exist only on the master node and be owned by root and have permissions 0400 (read only by root).
- *--make-shared-key* Generate a new random shared key. The key is a 256 random number encoded in base64.

Appendix D. Changes to Rocks Security Infrastructure

D.1. Rocks Password Infrastructure

D.1.1. Current Design

Currently, the root password for the frontend is elicited from the user during the frontend installation. Once the password is obtained, 3 crypted versions are generated from the root password. One version is the UNIX crypt version, which forms the root password for the cluster, another is the portable password format which is used for wordpress access. The third is an SHA version which is used as the MySQL root password. These three crypted passwords are stored in the database as shadow attributes.

When compute nodes are installed the root password that is stored in the database is transferred over to the compute node as an XML entity. This way we maintain root password consistency between client nodes and frontend. Any changes made to the root password on the frontend using the **passwd** command is propagated to the compute nodes using the 411 system.

D.1.2. Changes to the Design

From the above description of the existing design, we can see that the root password is stored in multiple places and used for multiple purposes. We determined that this was not the optimal security situation, and could use changes.

The way we've implemented changes are as follows:

1. As soon as the root password is obtained from the admin, it is converted to only the UNIX crypt format. It is used during the installation to set the root password, and is not used for any other purposes, and is not stored in the database.
2. Random passwords are generated for root access to the database and access to the wordpress engine.
3. The root account on each compute node gets a random root password.
4. The 411 system uses plugins to filter the content of the password and shadow files so that it does not transmit information about any accounts with UID under 500. Therefore, any changes made to the root password on the frontend is not propagated to the compute nodes.
5. To change the root password for the compute nodes, the admin may now use the "secure attribute" infrastructure in Rocks. The admin may also use the **passwd** command on the compute nodes to change the password without fear of override.



You may NOT use both the secure attributes and the **passwd** command to change the password on the compute nodes. If you do, the value in the secure attributes table always wins.

6. To change the root password on the frontend, you can now use the **passwd** command, and not worry about changing any of the other passwords in the cluster.

D.1.3. Changing the root password on a compute node

This section describes how to change the root password for compute nodes. It uses the rocks secure attribute infrastructure

```
# rocks set host sec_attr compute attr=root_pw
  Enter root_pw:
Confirm root_pw:
# rocks list host sec_attr compute
HOST      ATTR      VALUE                                     ENC
compute-0-0: root_pw $1$G55uZRrO$5MI7Nv14U5yWfKAlauqPE0 crypt
compute-0-1: root_pw $1$G55uZRrO$5MI7Nv14U5yWfKAlauqPE0 crypt
# rocks sync host sec_attr compute
```

1. **rocks set host sec_attr** creates the entry in the database.
2. **rocks sync host sec_attr** changes the root password on the compute nodes.



The admin can also set different root password for each node

D.2. Rocks Secure Attribute Infrastructure

As a replacement for shadow attributes, we introduced the Rocks secure attributes infrastructure. The reason for the change was internal, as shadow attributes were still accessible by the apache user. Secure attributes are not accessible by anyone but the root user.

The secure attributes are not synced along with any of the other regular attributes. The secure attributes are synced between hosts completely out of band using standard SSH.

D.2.1. Structure

The secure attributes infrastructure consists of two parts.

1. The secure attribute itself, which is stored in the database.
2. Plugins, in the form of python code, which defines how to act on the secure attribute. These plugins reside in `/opt/rocks/var/plugins/sec_attr/`. The plugins have a very simple API.

D.2.2. Plugins API

The Rocks secure attribute plugin class inherits a single class called `rocks.commands.sec_attr_plugin`.

This example is a very simple fictitious use-case. A service called **fake_service** runs on all compute nodes, and stores an SHA password in its configuration file. The configuration file for this service is called `/etc/fake_service.conf`. These are the contents of the config file.

```
# cat /etc/fake_service.conf
password=150b95f90c06f127a040a40f98582231369b6fda
```

D.2.2.1. Setting the Secure Attribute

This password can be stored in the database as a secure attribute. To store the password as a secure attribute, run

```
# rocks set host sec_attr compute attr=fake_svc_pw enc=sha
  Enter fake_svc_pw:
Confirm fake_svc_pw:
```

Enter the password string and confirm it (the same way you would when changing the root password).

D.2.2.2. Creating the plugin

Secure attribute plugins are located in the `/opt/rocks/var/plugins/sec_attr` directory. These plugins are written in python, and contain very simple API. There are two functions that must be a part of all plugins.

1. `get_sec_attr`: This function returns a the name of the secure attribute as stored in the database. In the example, this would return `fake_svc_pw`.



This name should match exactly the attribute name in the database. If multiple plugins return the same value for this function, then the last match wins.

2. `filter`: This function is the workhorse of the plugin. This takes in a single value , and performs the desired manipulation, filtering, and storage on it.

The plugin for the above example is `/opt/rocks/var/plugins/sec_attr/fake.py`. The name of the file is irrelevant. The command iterates over all available plugins and uses the last match on the return value of `get_sec_attr` function.

Example D-1. Plugin code for `fake_svc_pw` - `fake.py`

```
import rocks.commands
import os, sys, string

class plugin(rocks.commands.sec_attr_plugin):
    def get_sec_attr(self):
        return 'fake_svc_pw'

    def filter(self, value):
        f = open('/etc/fake_service.conf', 'w')
        f.write('password=%s\n' % value)
        f.close()
```

D.2.2.3. Syncing the attribute

Once the plugin has been created and the attribute has been assigned in the database, you can sync the plugin using the command

```
# rocks sync host sec_attr
# rocks run host compute command="cat /etc/fake_service.conf" collate=yes
compute-0-1: password=150b95f90c06f127a040a40f98582231369b6fda
compute-0-0: password=150b95f90c06f127a040a40f98582231369b6fda
```

Appendix E. Kickstart Nodes Reference

E.1. Rocks Base Nodes

E.1.1. 411-base

The packages and other common elements of the 411 Secure Information Service.

Parent Nodes:

- base

E.1.2. 411-client

Sets up the 411 Secure Information Service for clients. The 411 service will automatically configure itself when a file is published. Also puts all current 411 files from the frontend into the kickstart file for services that cannot tolerate a single 411 failure. Note that 411 can never guarantee full absolute success at any single time. It only offers consistency over the long term.

Parent Nodes:

- client

E.1.3. 411-server

Sets up the 411 Secure Information Service for Master nodes. Creates the RSA public and private keys for the cluster, and configures Apache for 411.

Parent Nodes:

- server

E.1.4. apache

Apache HTTP Server

Parent Nodes:

- server

E.1.5. auto-kickstart

Force a kickstart on power fail or system crash.

Parent Nodes:

- base

E.1.6. autofs

AutoFS for automounting home directories over NFS or the loopback device.

Parent Nodes:

- base

E.1.7. autofs-client

AutoFS Client

Parent Nodes:

- client

E.1.8. autofs-server

AutoFS server

Parent Nodes:

- server

E.1.9. base

Base class for all Rocks nodes. This should include compute nodes, frontend nodes, standalone laptops, computer labs, graphics nodes, nfs servers To achieve this level of flexibility this base class should have edges only to those classes that implement the core of Rocks.

Parent Nodes:

- client
- server

Children Nodes:

- 411-base
- auto-kickstart
- autofs
- custom-partition
- disk-stamp
- firewall
- foundation-base
- grub
- i18n
- ip-diag
- ipmi-base
- logrotate
- mail
- md-rebuild
- mosh
- node
- ntp
- postshell
- rpc
- scripting
- security
- selinux
- snmp-server
- syslog
- upstart
- yum
- yum-extras

E.1.10. bootstrap-dummy

Dummy terminal node for bootstrap-packages subgraph

Parent Nodes:

- bootstrap-packages-base

E.1.11. bootstrap-packages-base

Additional packages that may be needed on a server that was not originally built with Rocks. Tested on Native CentOS installed build machine.

Parent Nodes:

- devel-appliance

Children Nodes:

- bootstrap-dummy

E.1.12. ca

Sets up keys for the frontend Certificate Authority

Parent Nodes:

- security-server

E.1.13. cdr

CDR Tools (burnings, iso, ripping, mp3 encoding)

Parent Nodes:

- devel

E.1.14. central

A Rocks Cluster Central server. Can kickstart other servers over the network. Also contains general Apache security setup.

Parent Nodes:

- security-server

E.1.15. client

The 'client node' in the graph. This file is used as a connection point for other XML configuration nodes.

Parent Nodes:

- compute-appliance
- devel-appliance
- login
- nas

Children Nodes:

- 411-client
- autofs-client
- base
- client-firewall
- client-rootpw
- core
- dhcp-client
- emacs
- fail2ban
- foundation-client
- grub-client
- grub2-client
- ipmi-client
- mail-client
- networking-client
- node-certs
- ntp-client
- partitions-save
- pxeboot
- resolv
- routes-client
- rsh
- security-client
- serf-client
- serf-common
- shared-memory
- syslog-client
- x11

E.1.16. client-firewall

Client node firewall rules

Parent Nodes:

- client

E.1.17. client-rootpw

Parent Nodes:

- client

E.1.18. compilers

Parent Nodes:

- devel

Children Nodes:

- c-development
- fortran-development

E.1.19. compute

Compute Appliance Node. You should create your own extend-compute.xml file to add functionality to this appliance type.

Children Nodes:

- compute-appliance

E.1.20. compute-appliance

Rocks Cluster Compute Node.

Parent Nodes:

- compute

Children Nodes:

- client

E.1.21. custom-partition

Parent Nodes:

- base

Children Nodes:

- partition

E.1.22. devel

Parent Nodes:

- devel-appliance
- server

Children Nodes:

- cdr
- compilers
- docbook
- emacs

E.1.23. devel-server

Add Development Appliance Definition to the Frontend

Parent Nodes:

- devel-appliance
- server

E.1.24. dhcp-client

DHCP Client

Parent Nodes:

- client

E.1.25. dhcp-server

Setup the DHCP server for the cluster

Parent Nodes:

- server

E.1.26. disk-stamp

Take a root partition, and make it ours! This is the key to determining, on reinstalls, if we should save partitions (because the stamp is there) or blow away all the partitions on the disk (because the stamp isn't there).

Parent Nodes:

- base

E.1.27. dns-server

Configures a DNS nameserver for the cluster on the frontend. Both forward and reversed zones are defined using the database.

Parent Nodes:

- server

E.1.28. do-partition

Parent Nodes:

- partition

E.1.29. docbook

DOC Book support (needed to build rolls)

Parent Nodes:

- devel

E.1.30. emacs

Emacs OS

Parent Nodes:

- client
- devel

E.1.31. ethers-server-postshell

Intended to be run after postshell. Will re-source the probed-ethers.sh and then regenerate configs for the network interfaces, modprobe.conf. ethers-server.xml will have probed the interfaces to create probed-ethers.sh. A user can put "postshell" on the command line and then edit probed-ethers.sh to arbitrarily re-assign mac:interface definitions.

Parent Nodes:

- postshell

E.1.32. fail2ban

Install fail2ban on node with public connection

Parent Nodes:

- client
- server

E.1.33. firewall

Default firewall rules for all appliances. Disable everything but SSH.

Parent Nodes:

- base

E.1.34. fortran-development

Fortran

Parent Nodes:

- compilers

E.1.35. foundation-base

Rocks Foundation Packages

Parent Nodes:

- base

E.1.36. foundation-client

Rocks Foundation Packages

Parent Nodes:

- client

E.1.37. foundation-server

Rocks Foundation Packages

Parent Nodes:

- devel-appliance
- server

E.1.38. google-otp

Google OTP Authentication Support

Parent Nodes:

- login
- server

E.1.39. google-otp-server

Google OTP Authentication Support Server Side

Parent Nodes:

- server

E.1.40. grub

Boot loader support (GRand Unified Bootloader)

Parent Nodes:

- base

E.1.41. grub-client

Boot loader configuration for cluster appliances.

Parent Nodes:

- client

E.1.42. grub-server

Boot loader configuration for cluster frontends. Puts enough information on the command line to contact the central server on a reinstall.

Parent Nodes:

- server

E.1.43. grub2-client

Boot loader configuration for cluster appliances. Grub2 edition

Parent Nodes:

- client

E.1.44. homepage

Create a simple cluster home page

Parent Nodes:

- server

E.1.45. i18n

Internationalization support

Parent Nodes:

- base

E.1.46. insert-rolls

Populates the frontend database with initially selected rolls.

Parent Nodes:

- server

E.1.47. install

Do everything needed to kickstart compute nodes or, generally speaking, everything needed to kickstart any node from this machine.

Parent Nodes:

- server

E.1.48. ip-diag

TCP/IP Network diagnostic tools.

Parent Nodes:

- base

E.1.49. ipmi-base

Common IPMI Support

Parent Nodes:

- base

E.1.50. ipmi-client

Client Nodes

Parent Nodes:

- client

E.1.51. ipmi-server

Server Nodes

Parent Nodes:

- server

E.1.52. logrotate

Append rules to logrotate to prune files in /var/log

Parent Nodes:

- base

E.1.53. mail

Mail support for the cluster.

Parent Nodes:

- base

E.1.54. mail-client

Mail support for nodes in the cluster. These nodes do not accept mail for local delivery. They pass everything to a mail hub (generally the frontend).

Parent Nodes:

- client

E.1.55. mail-server

Mail hub for the cluster. This node serves as the post-office for the whole cluster.

Parent Nodes:

- server

E.1.56. md-rebuild

Moderate MD Raid Rebuild Speed

Parent Nodes:

- base

E.1.57. mosh

MOSH - Mobile Shell from MIT

Parent Nodes:

- base

E.1.58. nas

NAS Appliance Node. You should create your own extend-nas.xml file to add functionality to this appliance type.

Children Nodes:

- client
- nas-appliance
- nas-exports
- nas-partition

E.1.59. nas-appliance

Rocks Cluster NAS Node.

Parent Nodes:

- nas

E.1.60. nas-exports

Rocks Cluster NAS Node export file system definitions

Parent Nodes:

- nas

E.1.61. nas-partition

Define a partitioning layout for NAS nodes.

Parent Nodes:

- nas

E.1.62. networking-client

Configure a node's network

Parent Nodes:

- client

E.1.63. networking-server

Configure a node's network

Parent Nodes:

- server

E.1.64. nfs-server

NFS Servers for /export

Parent Nodes:

- server

E.1.65. node-certs

Place the root ssh public key from the installation machine on to the installed machine While we're at it suck over the SSL CA Cert/Key. This is a quick hack, we need a secure method for this to put our pants back on.

Parent Nodes:

- client

E.1.66. ntp

Network Time Protocol

Parent Nodes:

- base

E.1.67. ntp-client

Network Time Protocol

Parent Nodes:

- client

E.1.68. ntp-server

Network Time Protocol

Parent Nodes:

- server

E.1.69. null

An empty XML file.

Parent Nodes:

- wan

E.1.70. partition

Parent Nodes:

- custom-partition

Children Nodes:

- do-partition

E.1.71. partitions-save

Save the partition info from this node by sending a message back to the frontend. The frontend will store the partition info in the database

Parent Nodes:

- client

E.1.72. partitions-save-server

Save the partition info for a server node by sending a message to the local host. This will store the partition info in the database. This file is a copy of 'partitions-save' which is connected to the 'client' graph. We need two files because of ordering: we want to save the partitions as early as possible on client nodes so we order 'partitions-save' off 'HEAD', but for server nodes, we need to wait for the database to be configured, so we order 'partitions-save-server' off 'TAIL'.

Parent Nodes:

- server

E.1.73. postshell

Allow interaction after all post sections have run If postshell is on the command line, wait until /mnt/sysimage/tmp/postshell is removed interactively.

Parent Nodes:

- base

Children Nodes:

- ethers-server-postshell

E.1.74. profile-server

profile.d entries for servers in Rocks 7

Parent Nodes:

- server

E.1.75. pxe

Provides pre-execution environment (PXE) support for frontend machines.

Parent Nodes:

- server

E.1.76. pxeboot

PXE first configuration for client nodes

Parent Nodes:

- client

E.1.77. resolv

Generates the resolv.conf file for slave nodes. We point to all nameservers listed in the PrivateDNSServers entry in the app_globals table in the database. This node used to be named copy-resolv.xml.

Parent Nodes:

- client

E.1.78. rlo

Remote lights Out

Parent Nodes:

- server

E.1.79. rocks-devel

Rocks Development

Parent Nodes:

- devel-appliance
- server

E.1.80. rocks-dist

Distribution building with rocks-dist

Parent Nodes:

- devel-appliance
- server

E.1.81. root

The root node that is used by rolls in order to tell 'kpp' which node should be the root of the traversed graph

Children Nodes:

- server

E.1.82. routes-client

Set up static routes

Parent Nodes:

- client

E.1.83. routes-server

Set up static routes on server

Parent Nodes:

- server

E.1.84. rpc

RPC support

Parent Nodes:

- base

E.1.85. rsh

Enable RSH

Parent Nodes:

- client

E.1.86. screen-clusterinfo

Parent Nodes:

- server

E.1.87. screen-network-misc

Parent Nodes:

- server

E.1.88. screen-network-private

Parent Nodes:

- server

E.1.89. screen-network-public

Parent Nodes:

- server

E.1.90. screen-partitioning

Parent Nodes:

- server

E.1.91. screen-root-password

Parent Nodes:

- server

E.1.92. screen-timezone

Parent Nodes:

- server

E.1.93. scripting

Parent Nodes:

- base

Children Nodes:

- perl-development
- perl-modules
- python-development
- tcl-development

E.1.94. security

Security parent node for common security elements.

Parent Nodes:

- base

Children Nodes:

- ssh
- ssl

E.1.95. selinux

Configure SELinux

Parent Nodes:

- base

E.1.96. serf-client

serf

Parent Nodes:

- client

E.1.97. serf-common

serf

Parent Nodes:

- client
- server

E.1.98. serf-server

serf

Parent Nodes:

- server

E.1.99. server-firewall

Server node firewall rules

Parent Nodes:

- server

E.1.100. server-rootpw

\$Id: server-rootpw.xml,v 1.2 2011/06/21 23:26:04 anoop Exp \$ Root Password for the frontend

Parent Nodes:

- server

E.1.101. shared-memory

Configure the maximum size of a shared memory segment to be 3/4 of main memory

Parent Nodes:

- client

E.1.102. snmp-client

SNMP Client machine.

Parent Nodes:

- devel-appliance

- server

E.1.103. snmp-server

SNMP Server. Open all the tables to the public and customize a couple site variables.

Parent Nodes:

- base

E.1.104. ssh

Enable SSH

Parent Nodes:

- security

E.1.105. ssh-client

SSH Config for compute nodes and other non-frontend appliances. We are using one key pair among all SSH servers in the cluster. This implies we do not care about Man-in-the-Middle attacks. We have subverted the protection for these attacks for several releases (broadcastSSH). This logic should not be in the ssh.xml node so the frontend will generate its own keypair.

Parent Nodes:

- security-client

E.1.106. ssl

Open SSL support

Parent Nodes:

- security

E.1.107. ssl-client

Secure Sockets Layer for non-frontend nodes.

Parent Nodes:

- security-client

E.1.108. ssl-server

Secure Sockets Layer for frontend nodes.

Parent Nodes:

- security-server

E.1.109. syslog

Setup Syslog

Parent Nodes:

- base

E.1.110. syslog-client

Setup Syslog for client machine to forward messages

Parent Nodes:

- client

E.1.111. syslog-server

Setup Syslog for server to accept forwarded messages

Parent Nodes:

- server

E.1.112. tcl-development

Tcl support

Parent Nodes:

- scripting

E.1.113. upstart

Defines Upstart-compatible Services

Parent Nodes:

- base

E.1.114. wan

A new server installed over the wide area network, from a central.

Children Nodes:

- null

E.1.115. x11

X11 Desktop applications.

Parent Nodes:

- client
- devel-appliance
- server

E.1.116. x11-server

X11 on Server Nodes for CentOS 7

Parent Nodes:

- server

E.1.117. yum

Configure yum

Parent Nodes:

- base

E.1.118. yum-extras

Extra Yum Repositories. Yum.xml disables by default

Parent Nodes:

- base

E.1.119. bootstrap-packages-core

Node 'bootstrap-packages-core' referenced but not defined in this roll

E.1.120. c-development

Node 'c-development' referenced but not defined in this roll

E.1.121. core

Node 'core' referenced but not defined in this roll

E.1.122. database

Node 'database' referenced but not defined in this roll

E.1.123. database-data

Node 'database-data' referenced but not defined in this roll

E.1.124. database-schema

Node 'database-schema' referenced but not defined in this roll

E.1.125. database-sec

Node 'database-sec' referenced but not defined in this roll

E.1.126. devel-appliance

Node 'devel-appliance' referenced but not defined in this roll

E.1.127. login

Node 'login' referenced but not defined in this roll

E.1.128. login-server

Node 'login-server' referenced but not defined in this roll

E.1.129. node

Node 'node' referenced but not defined in this roll

E.1.130. perl-development

Node 'perl-development' referenced but not defined in this roll

E.1.131. perl-modules

Node 'perl-modules' referenced but not defined in this roll

E.1.132. python-development

Node 'python-development' referenced but not defined in this roll

E.1.133. security-client

Node 'security-client' referenced but not defined in this roll

E.1.134. security-server

Node 'security-server' referenced but not defined in this roll

E.1.135. server

Node 'server' referenced but not defined in this roll

Appendix F. Rocks® Copyright and Trademark

F.1. Copyright Statement

Rocks(r)
www.rocksclusters.org
version 6.2 (SideWinder)
version 7.0 (Manzanita)

Copyright (c) 2000 - 2017 The Regents of the University of California.
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice unmodified and in its entirety, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising and press materials, printed or electronic, mentioning features or use of this software must display the following acknowledgement:

"This product includes software developed by the Rocks(r)
Cluster Group at the San Diego Supercomputer Center at the
University of California, San Diego and its contributors."

4. Except as permitted for the purposes of acknowledgment in paragraph 3, neither the name or logo of this software nor the names of its authors may be used to endorse or promote products derived from this software without specific prior written permission. The name of the software includes the following terms, and any derivatives thereof: "Rocks", "Rocks Clusters", and "Avalanche Installer". For licensing of the associated name, interested parties should contact Technology Transfer & Intellectual Property Services, University of California, San Diego, 9500 Gilman Drive, Mail Code 0910, La Jolla, CA 92093-0910, Ph: (858) 534-5815, FAX: (858) 534-7345, E-MAIL:invent@ucsd.edu

THIS SOFTWARE IS PROVIDED BY THE REGENTS AND CONTRIBUTORS "AS IS AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE

OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN
IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

F.2. Trademark Licensing

The Rocks® logo (Figure F-1) and the names Rocks®, Rocks Clusters™, and Avalanche Installer™ have been submitted to the United States Patent and Trademark Office to become registered trademarks.

Commercial entities that wish to use any of the above names or logo in a product name or marketing material, are required to get written permission from the Technology Transfer and Intellectual Property Services Office at UCSD <invent@ucsd.edu>.

Figure F-1. Rocks® logo



Appendix G. Common Licenses

G.1. Artistic License

The "Artistic License"

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

Definitions:

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:

a) place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b) use the modified Package only within your corporation or organization.

c) rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a) distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b) accompany the distribution with the machine-readable source of the Package with your modifications.

c) give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d) make other distribution arrangements with the Copyright Holder.

5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Package's interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.

6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.

7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to

emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.

8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Package's interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.

9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

G.2. Apache v2.0

Apache License
Version 2.0, January 2004
<http://www.apache.org/licenses/>

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical

transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
 - (a) You must give any other recipients of the Work or Derivative Works a copy of this License; and
 - (b) You must cause any modified files to carry prominent notices stating that You changed the files; and
 - (c) You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and
 - (d) If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. Trademarks. This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each

Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

G.3. GNU General Public License v1

GNU GENERAL PUBLIC LICENSE
Version 1, February 1989

Copyright (C) 1989 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307, USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The license agreements of most software companies try to keep users at the mercy of those companies. By contrast, our General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. The General Public License applies to the Free Software Foundation's software and to any other program whose authors commit to using it. You can use it for your programs, too.

When we speak of free software, we are referring to freedom, not price. Specifically, the General Public License is designed to make sure that you have the freedom to give away or sell copies of free software, that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of a such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must tell them their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any work containing the Program or a portion of it, either verbatim or with modifications. Each licensee is addressed as "you".

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this General Public License and to the absence of any warranty; and give any other recipients of the Program a copy of this General Public License along with the Program. You may charge a fee for the physical act of transferring a copy.

2. You may modify your copy or copies of the Program or any portion of it, and copy and distribute such modifications under the terms of Paragraph 1 above, provided that you also do the following:

a) cause the modified files to carry prominent notices stating that

you changed the files and the date of any change; and

b) cause the whole of any work that you distribute or publish, that in whole or in part contains the Program or any part thereof, either with or without modifications, to be licensed at no charge to all third parties under the terms of this General Public License (except that you may choose to grant warranty protection to some or all third parties, at your option).

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the simplest and most usual way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this General Public License.

d) You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

Mere aggregation of another independent work with the Program (or its derivative) on a volume of a storage or distribution medium does not bring the other work under the scope of these terms.

3. You may copy and distribute the Program (or a portion or derivative of it, under Paragraph 2) in object code or executable form under the terms of Paragraphs 1 and 2 above provided that you also do one of the following:

a) accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Paragraphs 1 and 2 above; or,

b) accompany it with a written offer, valid for at least three years, to give any third party free (except for a nominal charge for the cost of distribution) a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Paragraphs 1 and 2 above; or,

c) accompany it with the information you received as to where the corresponding source code may be obtained. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form alone.)

Source code for a work means the preferred form of the work for making modifications to it. For an executable file, complete source code means all the source code for all modules it contains; but, as a special exception, it need not include source code for modules which are standard libraries that accompany the operating system on which the executable file runs, or for standard header files or definitions files that accompany that operating system.

4. You may not copy, modify, sublicense, distribute or transfer the Program except as expressly provided under this General Public License. Any attempt otherwise to copy, modify, sublicense, distribute or transfer the Program is void, and will automatically terminate your rights to use

the Program under this License. However, parties who have received copies, or rights to use copies, from you under this General Public License will not have their licenses terminated so long as such parties remain in full compliance.

5. By copying, distributing or modifying the Program (or any work based on the Program) you indicate your acceptance of this license to do so, and all its terms and conditions.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein.

7. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of the license which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of the license, you may choose any version ever published by the Free Software Foundation.

8. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

9. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

10. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to humanity, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) 19yy <name of author>
```

```
This program is free software; you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation; either version 1, or (at your option)  
any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program; if not, write to the Free Software Foundation,  
Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307, USA.
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) 19xx name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the  
program 'Gnomovision' (a program to direct compilers to make passes  
at assemblers) written by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989
```

Ty Coon, President of Vice

That's all there is to it!

G.4. GNU General Public License v2

GNU GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so

that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
- b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
- c) If the modified program normally reads commands interactively

when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:

- a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
- c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include

anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the

integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED

TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the program's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This program is free software; you can redistribute it and/or modify  
it under the terms of the GNU General Public License as published by  
the Free Software Foundation; either version 2 of the License, or  
(at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the  
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License  
along with this program; if not, write to the Free Software  
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
```

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author  
Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'.  
This is free software, and you are welcome to redistribute it  
under certain conditions; type `show c' for details.
```

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program  
'Gnomovision' (which makes passes at compilers) written by James Hacker.
```

<signature of Ty Coon>, 1 April 1989
Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

G.5. GNU Lesser General Public License v2.1

GNU LESSER GENERAL PUBLIC LICENSE
Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts
as the successor of the GNU Library Public License, version 2, hence
the version number 2.1.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis

or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original author's reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the user's freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free

programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is less protective of the users' freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

GNU LESSER GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any

warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.
- d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public

License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6,

whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)

b) Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the user's computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.

c) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.

d) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.

e) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.

b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by

all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR

OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or  
modify it under the terms of the GNU Lesser General Public  
License as published by the Free Software Foundation; either  
version 2.1 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU  
Lesser General Public License for more details.
```

```
You should have received a copy of the GNU Lesser General Public  
License along with this library; if not, write to the Free Software  
Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the
```

library 'Frob' (a library for tweaking knobs) written by James Random Hacker.

<signature of Ty Coon>, 1 April 1990
Ty Coon, President of Vice

That's all there is to it!

G.6. GNU Library General Public License v2

GNU LIBRARY GENERAL PUBLIC LICENSE
Version 2, June 1991

Copyright (C) 1991 Free Software Foundation, Inc.
675 Mass Ave, Cambridge, MA 02139, USA
Everyone is permitted to copy and distribute verbatim copies
of this license document, but changing it is not allowed.

[This is the first released version of the library GPL. It is
numbered 2 because it goes with version 2 of the ordinary GPL.]

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Library General Public License, applies to some specially designated Free Software Foundation software, and to any other libraries whose authors decide to use it. You can use it for your libraries, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library, or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link a program with the library, you must provide complete object files to the recipients so that they can relink them with the library, after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

Our method of protecting your rights has two steps: (1) copyright the library, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the library.

Also, for each distributor's protection, we want to make certain that everyone understands that there is no warranty for this free library. If the library is modified by someone else and passed on, we want its recipients to know that what they have is not the original version, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that companies distributing free software will individually obtain patent licenses, thus in effect transforming the program into proprietary software. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License, which was designed for utility programs. This license, the GNU Library General Public License, applies to certain designated libraries. This license is quite different from the ordinary one; be sure to read it in full, and don't assume that anything in it is the same as in the ordinary license.

The reason we have a separate public license for some libraries is that they blur the distinction we usually make between modifying or adding to a program and simply using it. Linking a program with a library, without changing the library, is in some sense simply using the library, and is analogous to running a utility program or application program. However, in a textual and legal sense, the linked executable is a combined work, a derivative of the original library, and the ordinary General Public License treats it as such.

Because of this blurred distinction, using the ordinary General Public License for libraries did not effectively promote software sharing, because most developers did not use the libraries. We concluded that weaker conditions might promote sharing better.

However, unrestricted linking of non-free programs would deprive the users of those programs of all benefit from the free status of the libraries themselves. This Library General Public License is intended to permit developers of non-free programs to use free libraries, while preserving your freedom as a user of such programs to change the free libraries that are incorporated in them. (We have not seen how to achieve this as regards changes in header files, but we have achieved it as regards changes in the actual functions of the Library.) The hope is that this will lead to faster development of free libraries.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, while the latter only works together with the library.

Note that it is possible for a library to be covered by the ordinary General Public License rather than by this special one.

GNU LIBRARY GENERAL PUBLIC LICENSE
TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Library General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

- a) The modified work must itself be a software library.
- b) You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.

c) You must cause the whole of the work to be licensed at no charge to all third parties under the terms of this License.

d) If a facility in the modified Library refers to a function or a table of data to be supplied by an application program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library.

In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or

derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also compile or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customer's own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a) Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b) Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- c) If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- d) Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:

- a) Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
- b) Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

8. You may not copy, modify, sublicense, link with, or distribute

the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.

10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.

11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply, and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the

original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

13. The Free Software Foundation may publish revised and/or new versions of the Library General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.

14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Libraries

If you develop a new library, and you want it to be of the greatest possible use to the public, we recommend making it free software that

everyone can redistribute and change. You can do so by permitting redistribution under these terms (or, alternatively, under the terms of the ordinary General Public License).

To apply these terms, attach the following notices to the library. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

```
<one line to give the library's name and a brief idea of what it does.>  
Copyright (C) <year> <name of author>
```

```
This library is free software; you can redistribute it and/or  
modify it under the terms of the GNU Library General Public  
License as published by the Free Software Foundation; either  
version 2 of the License, or (at your option) any later version.
```

```
This library is distributed in the hope that it will be useful,  
but WITHOUT ANY WARRANTY; without even the implied warranty of  
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU  
Library General Public License for more details.
```

```
You should have received a copy of the GNU Library General Public  
License along with this library; if not, write to the Free  
Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.
```

Also add information on how to contact you by electronic and paper mail.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the library, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the  
library 'Frob' (a library for tweaking knobs) written by James Random Hacker.
```

```
<signature of Ty Coon>, 1 April 1990  
Ty Coon, President of Vice
```

That's all there is to it!

G.7. Python Software Foundation License v2

PYTHON SOFTWARE FOUNDATION LICENSE VERSION 2

1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee") accessing and otherwise using this software ("Python") in source or binary form and its associated documentation.

2. Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide

license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python alone or in any derivative version, provided, however, that PSF's License Agreement and PSF's notice of copyright, i.e., "Copyright (c) 2001, 2002, 2003, 2004 Python Software Foundation; All Rights Reserved" are retained in Python alone or in any derivative version prepared by Licensee.

3. In the event Licensee prepares a derivative work that is based on or incorporates Python or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python.

4. PSF is making Python available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.

5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.

7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.

8. By copying, installing or otherwise using Python, Licensee agrees to be bound by the terms and conditions of this License Agreement.

Appendix H. Package 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@ GIT repository¹.

H.1. anaconda

GNU General Public License v2

H.2. ant

Apache v2.0

H.2.1. dom

This license came from:

<http://www.w3.org/Consortium/Legal/copyright-software-19980720>

W3C SOFTWARE NOTICE AND LICENSE

Copyright © 1994-2002 World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>

This W3C work (including software, documents, or other related items) is being provided by the copyright holders under the following license. By obtaining, using and/or copying this work, you (the licensee) agree that you have read, understood, and will comply with the following terms and conditions:

Permission to use, copy, modify, and distribute this software and its documentation, with or without modification, for any purpose and without fee or royalty is hereby granted, provided that you include the following on ALL copies of the software and documentation or portions thereof, including modifications, that you make:

1. The full text of this NOTICE in a location viewable to users of the redistributed or derivative work.
2. Any pre-existing intellectual property disclaimers, notices, or terms and conditions. If none exist, a short notice of the following form (hypertext is preferred, text is permitted) should be used within the body of any redistributed or derivative code: "Copyright © [date-of-software] World Wide Web Consortium, (Massachusetts Institute of Technology, Institut National de Recherche en Informatique et en Automatique, Keio University). All Rights Reserved. <http://www.w3.org/Consortium/Legal/>"
3. Notice of any changes or modifications to the W3C files, including the date changes were made. (We recommend you provide URIs to the location from which the code is derived.)

THIS SOFTWARE AND DOCUMENTATION IS PROVIDED "AS IS," AND COPYRIGHT HOLDERS MAKE

NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE OR DOCUMENTATION WILL NOT INFRINGE ANY THIRD PARTY PATENTS, COPYRIGHTS, TRADEMARKS OR OTHER RIGHTS.

COPYRIGHT HOLDERS WILL NOT BE LIABLE FOR ANY DIRECT, INDIRECT, SPECIAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF ANY USE OF THE SOFTWARE OR DOCUMENTATION.

The name and trademarks of copyright holders may NOT be used in advertising or publicity pertaining to the software without specific, written prior permission. Title to copyright in this software and any associated documentation will at all times remain with copyright holders.

H.2.2. sax

This license came from: <http://www.megginson.com/SAX/copying.html>

However please note future versions of SAX may be covered under <http://saxproject.org/?selected=pd>

This page is now out of date -- see the new SAX site at <http://www.saxproject.org/> for more up-to-date releases and other information. Please change your bookmarks.

SAX2 is Free!

I hereby abandon any property rights to SAX 2.0 (the Simple API for XML), and release all of the SAX 2.0 source code, compiled code, and documentation contained in this distribution into the Public Domain. SAX comes with NO WARRANTY or guarantee of fitness for any purpose.

David Megginson, david@megginson.com
2000-05-05

H.2.3. xerces

The Apache Software License, Version 1.1

Copyright (c) 1999-2002 The Apache Software Foundation. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright

notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The end-user documentation included with the redistribution, if any, must include the following acknowledgment:
"This product includes software developed by the Apache Software Foundation (<http://www.apache.org/>)."
Alternately, this acknowledgment may appear in the software itself, if and wherever such third-party acknowledgments normally appear.
4. The names "Xerces" and "Apache Software Foundation" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact apache@apache.org.
5. Products derived from this software may not be called "Apache", nor may "Apache" appear in their name, without prior written permission of the Apache Software Foundation.

THIS SOFTWARE IS PROVIDED "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE APACHE SOFTWARE FOUNDATION OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

=====

This software consists of voluntary contributions made by many individuals on behalf of the Apache Software Foundation and was originally based on software copyright (c) 1999, International Business Machines, Inc., <http://www.ibm.com>. For more information on the Apache Software Foundation, please see www.apache.org

H.3. coreutils

GNU General Public License v2

H.4. cvs

GNU General Public License v1

H.5. eGenix mx

EGENIX.COM PUBLIC LICENSE AGREEMENT

VERSION 1.0.0

1. Introduction

This "License Agreement" is between eGenix.com Software, Skills and Services GmbH ("eGenix.com"), having an office at Pastor-Loeh-Str. 48, D-40764 Langenfeld, Germany, and the Individual or Organization ("Licensee") accessing and otherwise using this software in source or binary form and its associated documentation ("the Software").

2. License

Subject to the terms and conditions of this eGenix.com Public License Agreement, eGenix.com hereby grants Licensee a non-exclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use the Software alone or in any derivative version, provided, however, that the eGenix.com Public License Agreement is retained in the Software, or in any derivative version of the Software prepared by Licensee.

3. NO WARRANTY

eGenix.com is making the Software available to Licensee on an "AS IS" basis. SUBJECT TO ANY STATUTORY WARRANTIES WHICH CAN NOT BE EXCLUDED, EGENIX.COM MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, EGENIX.COM MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF THE SOFTWARE WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.

4. LIMITATION OF LIABILITY

EGENIX.COM SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF THE SOFTWARE FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF BUSINESS PROFITS, BUSINESS INTERRUPTION, LOSS OF BUSINESS INFORMATION, OR OTHER PECUNIARY LOSS) AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THE SOFTWARE, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.

SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THE ABOVE EXCLUSION OR LIMITATION MAY NOT APPLY TO LICENSEE.

5. Termination

This License Agreement will automatically terminate upon a material breach of its terms and conditions.

6. General

Nothing in this License Agreement affects any statutory rights of consumers that cannot be waived or limited by contract.

Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between eGenix.com and Licensee.

If any provision of this License Agreement shall be unlawful, void, or for any reason unenforceable, such provision shall be modified to the extent necessary to render it enforceable without losing its intent, or, if no such modification is possible, be severed from this License Agreement and shall not affect the validity and enforceability of the remaining provisions of this License Agreement.

This License Agreement shall be governed by and interpreted in all respects by the law of Germany, excluding conflict of law provisions. It shall not be governed by the United Nations Convention on Contracts for International Sale of Goods.

This License Agreement does not grant permission to use eGenix.com trademarks or trade names in a trademark sense to endorse or promote products or services of Licensee, or any third party.

The controlling language of this License Agreement is English. If Licensee has received a translation into another language, it has been provided for Licensee's convenience only.

14. Agreement

By downloading, copying, installing or otherwise using the Software, Licensee agrees to be bound by the terms and conditions of this License Agreement.

H.6. FireFox

MOZILLA PUBLIC LICENSE
Version 1.1

1. Definitions.

1.0.1. "Commercial Use" means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor" means each entity that creates or contributes to

the creation of Modifications.

1.2. "Contributor Version" means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifications made by that particular Contributor.

1.3. "Covered Code" means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism" means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable" means Covered Code in any form other than Source Code.

1.6. "Initial Developer" means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work" means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License" means this document.

1.8.1. "Licensable" means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications" means any addition to or deletion from the substance or structure of either the Original Code or any previous Modifications. When Covered Code is released as a series of files, a Modification is:

A. Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

B. Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code" means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims" means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code" means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available

for no charge.

1.12. "You" (or "Your") means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and

(b) under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

(c) the licenses granted in this Section 2.1(a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of

Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) the licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifications of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version available; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subsequent version of that particular Modification has been made available to such recipients. You are responsible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Developer and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Covered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims.

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowledge after the Modification is made available as described in Section 3.2, Contributor shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs.

If Contributor's Modifications include an application programming interface and Contributor has knowledge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the LEGAL file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4(a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipients' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. However, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Section 3.1-3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered

Code or ownership rights under a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code they affect. Such description must be included in the LEGAL file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such description must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions.

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions.

Once Covered Code has been published under a particular version of the License, You may always continue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works.

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape",

"MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Public License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Contributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. DISCLAIMER OF WARRANTY.

COVERED CODE IS PROVIDED UNDER THIS LICENSE ON AN "AS IS" BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED CODE IS FREE OF DEFECTS, MERCHANTABILITY, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED CODE IS WITH YOU. SHOULD ANY COVERED CODE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED CODE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

8. TERMINATION.

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

(a) such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Contributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not withdrawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

(b) any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that

Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contributor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (excluding distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED CODE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

10. U.S. GOVERNMENT END USERS.

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to disputes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including

without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. MULTIPLE-LICENSED CODE.

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the NPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

EXHIBIT A -Mozilla Public License.

"The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/MPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.
Portions created by _____ are Copyright (C) _____
_____. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[_____] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [_____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [_____] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [_____] License."

[NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should

use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.]

AMENDMENTS

The Netscape Public License Version 1.1 ("NPL") consists of the Mozilla Public License Version 1.1 with the following Amendments, including Exhibit A-Netscape Public License. Files identified with "Exhibit A-Netscape Public License" are governed by the Netscape Public License Version 1.1.

Additional Terms applicable to the Netscape Public License.

I. Effect.

These additional terms described in this Netscape Public License -- Amendments shall apply to the Mozilla Communicator client code and to all Covered Code under this License.

II. "Netscape's Branded Code" means Covered Code that Netscape distributes and/or permits others to distribute under one or more trademark(s) which are controlled by Netscape but which are not licensed for use under this License.

III. Netscape and logo.

This License does not grant any rights to use the trademarks "Netscape", the "Netscape N and horizon" logo or the "Netscape lighthouse" logo, "Netcenter", "Gecko", "Java" or "JavaScript", "Smart Browsing" even if such marks are included in the Original Code or Modifications.

IV. Inability to Comply Due to Contractual Obligation.

Prior to licensing the Original Code under this License, Netscape has licensed third party code for use in Netscape's Branded Code. To the extent that Netscape is limited contractually from making such third party code available under this License, Netscape may choose to reintegrate such code into Covered Code without being required to distribute such code in Source Code form, even if such code would otherwise be considered "Modifications" under this License.

V. Use of Modifications and Covered Code by Initial Developer.

V.1. In General.

The obligations of Section 3 apply to Netscape, except to the extent specified in this Amendment, Section V.2 and V.3.

V.2. Other Products.

Netscape may include Covered Code in products other than the Netscape's Branded Code which are released by Netscape during the two (2) years following the release date of the Original Code, without such additional products becoming subject to the terms of this License, and may license such additional products on different terms from those contained in this License.

V.3. Alternative Licensing.

Netscape may license the Source Code of Netscape's Branded

Code, including Modifications incorporated therein, without such Netscape Branded Code becoming subject to the terms of this License, and may license such Netscape Branded Code on different terms from those contained in this License.

VI. Litigation.

Notwithstanding the limitations of Section 11 above, the provisions regarding litigation in Section 11(a), (b) and (c) of the License shall apply to all disputes relating to this License.

EXHIBIT A-Netscape Public License.

"The contents of this file are subject to the Netscape Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at <http://www.mozilla.org/NPL/>

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations under the License.

The Original Code is Mozilla Communicator client code, released March 31, 1998.

The Initial Developer of the Original Code is Netscape Communications Corporation. Portions created by Netscape are Copyright (C) 1998-1999 Netscape Communications Corporation. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[____] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [____] License and not to allow others to use your version of this file under the NPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [____] License. If you do not delete the provisions above, a recipient may use your version of this file under either the NPL or the [____] License."

H.7. gawk

GNU General Public License v2

H.8. gd

Portions copyright 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002 by Cold Spring Harbor Laboratory. Funded under Grant P41-RR02188 by the National Institutes of Health.

Portions copyright 1996, 1997, 1998, 1999, 2000, 2001, 2002 by Boutell.Com, Inc.

Portions relating to GD2 format copyright 1999, 2000, 2001, 2002 Philip Warner.

Portions relating to PNG copyright 1999, 2000, 2001, 2002 Greg Roelofs.

Portions relating to gdttf.c copyright 1999, 2000, 2001, 2002 John Ellson (ellson@lucent.com).

Portions relating to gdft.c copyright 2001, 2002 John Ellson (ellson@lucent.com).

Portions relating to JPEG and to color quantization copyright 2000, 2001, 2002, Doug Becker and copyright (C) 1994, 1995, 1996, 1997, 1998, 1999, 2000, 2001, 2002, Thomas G. Lane. This software is based in part on the work of the Independent JPEG Group. See the file README-JPEG.TXT for more information.

Portions relating to WBMP copyright 2000, 2001, 2002 Maurice Szmurlo and Johan Van den Brande.

Permission has been granted to copy, distribute and modify gd in any context without fee, including a commercial application, provided that this notice is present in user-accessible supporting documentation.

This does not affect your ownership of the derived work itself, and the intent is to assure proper credit for the authors of gd, not to interfere with your productive use of gd. If you have questions, ask. "Derived works" includes all programs that utilize the library. Credit must be given in user-accessible documentation.

This software is provided "AS IS." The copyright holders disclaim all warranties, either express or implied, including but not limited to implied warranties of merchantability and fitness for a particular purpose, with respect to this code and accompanying documentation.

Although their code does not appear in gd, the authors wish to thank David Koblas, David Rowley, and Hutchison Avenue Software Corporation for their prior contributions.

H.9. graphviz

Common Public License Version 1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

1. DEFINITIONS

"Contribution" means:

a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and

b) in the case of each subsequent Contributor:

i) changes to the Program, and

ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations

which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

a) it complies with the terms and conditions of this Agreement; and

b) its license agreement:

i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

a) it must be made available under this Agreement; and

b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the

terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to software (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Contributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

H.10. kudzu

GNU General Public License v2

H.11. libxml2

Except where otherwise noted in the source code (e.g. the files hash.c, list.c and the trio files, which are covered by a similar licence but with different Copyright notices) all the files are:

Copyright (C) 1998-2003 Daniel Veillard. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE DANIEL VEILLARD BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name of Daniel Veillard shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization from him.

H.12. libxml2doc

GNU Lesser General Public License v2.1

Copyright (C) 2003, 2004, 2005 Paul Boddie <paul@boddie.org.uk>

This library is free software; you can redistribute it and/or modify it under the terms of the GNU Lesser General Public License as published by the Free Software Foundation; either version 2.1 of the License, or (at your option) any later version.

This library is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU Lesser General Public License for more details.

You should have received a copy of the GNU Lesser General Public License along with this library; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301, USA

H.13. mysql

GNU General Public License v2

H.13.1. FLOSS License Exception

MySQL FLOSS License Exception

The MySQL AB Exception for Free/Libre and Open Source Software-only Applications Using MySQL Client Libraries (the "FLOSS Exception").

Version 0.6, 7 March 2007

Exception Intent

We want specified Free/Libre and Open Source Software ("FLOSS") applications to be able to use specified GPL-licensed MySQL client libraries (the "Program") despite the fact that not all FLOSS licenses are compatible with version 2 of the GNU General Public License (the "GPL").

Legal Terms and Conditions

As a special exception to the terms and conditions of version 2.0 of the GPL:

1. You are free to distribute a Derivative Work that is formed entirely from the Program and one or more works (each, a "FLOSS Work") licensed under one or more of the licenses listed below in section 1, as long as:
 - a. You obey the GPL in all respects for the Program and the Derivative Work, except for identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves,
 - b. all identifiable sections of the Derivative Work which are not derived from the Program, and which can reasonably be considered independent and separate works in themselves,
 - i. are distributed subject to one of the FLOSS licenses listed below, and
 - ii. the object code or executable form of those sections are accompanied by the complete corresponding machine-readable source code for those sections on the same medium and under the same FLOSS license as the corresponding object code or executable forms of those sections, and
 - c. any works which are aggregated with the Program or with a Derivative Work on a volume of a storage or distribution medium in accordance with the GPL, can reasonably be considered independent and separate works in themselves which are not derivatives of either the Program, a Derivative Work or a FLOSS Work.

If the above conditions are not met, then the Program may only

be copied, modified, distributed or used under the terms and conditions of the GPL or another valid licensing option from MySQL AB.

2. FLOSS License List

License name Version(s)/Copyright Date
Academic Free License 2.0
Apache Software License 1.0/1.1/2.0
Apple Public Source License 2.0
Artistic license From Perl 5.8.0
BSD license "July 22 1999"
Common Development and Distribution License (CDDL) 1.0
Common Public License 1.0
Eclipse Public License 1.0
GNU Library or "Lesser" General Public License (LGPL) 2.0/2.1
Jabber Open Source License 1.0
MIT license (As listed in file MIT-License.txt) ---
Mozilla Public License (MPL) 1.0/1.1
Open Software License 2.0
OpenSSL license (with original SSLeay license) "2003" ("1998")
PHP License 3.0
Python license (CNRI Python License) ---
Python Software Foundation License 2.1.1
Sleepycat License "1999"
University of Illinois/NCSA Open Source License ---
W3C License "2001"
X11 License "2001"
Zlib/libpng License ---
Zope Public License 2.0

Due to the many variants of some of the above licenses, we require that any version follow the 2003 version of the Free Software Foundation's Free Software Definition (<http://www.gnu.org/philosophy/free-sw.html>) or version 1.9 of the Open Source Definition by the Open Source Initiative (<http://www.opensource.org/docs/definition.php>).

3. Definitions

- a. Terms used, but not defined, herein shall have the meaning provided in the GPL.
- b. Derivative Work means a derivative work under copyright law.

4. **Applicability:** This FLOSS Exception applies to all Programs that contain a notice placed by MySQL AB saying that the Program may be distributed under the terms of this FLOSS Exception. If you create or distribute a work which is a Derivative Work of both the Program and any other work licensed under the GPL, then this FLOSS Exception is not available for that work; thus, you must remove the FLOSS Exception notice from that work and comply with the GPL in all respects, including by retaining all GPL notices. You may choose to redistribute a copy of the Program exclusively under the terms of the GPL by removing the FLOSS Exception notice from that copy of the Program, provided that the copy has

never been modified by you or any third party.

Appendix A. Qualified Libraries and Packages

The following is a non-exhaustive list of libraries and packages which are covered by the FLOSS License Exception. Please note that this appendix is provided merely as an additional service to specific FLOSS projects wishing to simplify licensing information for their users. Compliance with one of the licenses noted under the "FLOSS license list" section remains a prerequisite.

Package Name	Qualifying License and Version
Apache Portable Runtime (APR)	Apache Software License 2.0

H.14. ncurses

Copyright (c) 1998-2003,2004 Free Software Foundation, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, distribute with modifications, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE ABOVE COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Except as contained in this notice, the name(s) of the above copyright holders shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization.

H.15. numarray

Copyright (C) 2001,2002,2003 Association of Universities for Research in Astronomy (AURA)

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name of AURA and its representatives may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY AURA "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL AURA BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

H.16. Numeric

Python Software Foundation License v2

This software was originally written by Jim Hugunin when he was at MIT, later at CNRI.

Then Paul Dubois of LLNL became the "keeper". LLNL has released any contributions to it by LLNL personnel for free redistribution. However, many people have contributed to Numpy. It is a classic open source effort that nobody "owns". So, Numerical Python is placed in the SourceForge under the Python license. The LLNL license is in Legal.htm.

-- Paul F. Dubois, Jan. 13, 2000.

H.16.1. LLNL License

*** Legal Notice for all LLNL-contributed files ***

Copyright (c) 1996. The Regents of the University of California. All rights reserved.

Permission to use, copy, modify, and distribute this software for any purpose without fee is hereby granted, provided that this entire notice is included in all copies of any software which is or includes a copy or modification of this software and in all copies of the supporting documentation for such software.

This work was produced at the University of California, Lawrence Livermore National Laboratory under contract no. W-7405-ENG-48 between the U.S. Department of Energy and The Regents of the University of California for the operation of UC LLNL.

DISCLAIMER

This software was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights. Reference herein to any specific commercial products, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

H.17. perl

GNU General Public License v1, GNU General Public License v2, Artistic License

Perl Kit, Version 5

Copyright (C) 1993, 1994, 1995, 1996, 1997, 1998, 1999,
2000, 2001, 2002, 2003, 2004, 2005, 2006 by Larry Wall and others

All rights reserved.

This program is free software; you can redistribute it and/or modify it under the terms of either:

- a) the GNU General Public License as published by the Free Software Foundation; either version 1, or (at your option) any later version, or
- b) the "Artistic License" which comes with this Kit.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See either the GNU General Public License or the Artistic License for more details.

You should have received a copy of the Artistic License with this

Kit, in the file named "Artistic". If not, I'll be glad to provide one.

You should also have received a copy of the GNU General Public License along with this program in the file named "Copying". If not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307, USA or visit their web page on the internet at <http://www.gnu.org/copyleft/gpl.html>.

For those of you that choose to use the GNU General Public License, my interpretation of the GNU General Public License is that no Perl script falls under the terms of the GPL unless you explicitly put said script under the terms of the GPL yourself. Furthermore, any object code linked with perl does not automatically fall under the terms of the GPL, provided such object code only adds definitions of subroutines and variables, and does not otherwise impair the resulting interpreter from executing any standard Perl script. I consider linking in C subroutines in this manner to be the moral equivalent of defining subroutines in the Perl language itself. You may sell such an object file as proprietary provided that you provide or offer to provide the Perl source, as specified by the GNU General Public License. (This is merely an alternate way of specifying input to the program.) You may also sell a binary produced by the dumping of a running Perl script that belongs to you, provided that you provide or offer to provide the Perl source as specified by the GPL. (The fact that a Perl interpreter and your code are in the same binary file is, in this case, a form of mere aggregation.) This is my interpretation of the GPL. If you still have concerns or difficulties understanding my intent, feel free to contact me. Of course, the Artistic License spells all this out for your protection, so you may prefer to use that.

H.18. perl tk

Perl License

Copyright (c) 1995-2003 Nick Ing-Simmons. All rights reserved.
This package is free software; you can redistribute it and/or modify it under the same terms as Perl itself, with the exception of the files in the pTk sub-directory which have separate terms derived from those of the original Tk4.0 sources and/or Tix.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

See pTk/license.terms for details of this Tk license,
and pTk/Tix.license for the Tix license.

H.18.1. pTk

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS AND DISTRIBUTORS SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT. THIS SOFTWARE IS PROVIDED ON AN "AS IS" BASIS, AND THE AUTHORS AND DISTRIBUTORS HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

GOVERNMENT USE: If you are acquiring this software on behalf of the U.S. government, the Government shall have only "Restricted Rights" in the software and related documentation as defined in the Federal Acquisition Regulations (FARs) in Clause 52.227.19 (c) (2). If you are acquiring the software on behalf of the Department of Defense, the software shall be classified as "Commercial Computer Software" and the Government shall have only "Restricted Rights" as defined in Clause 252.227-7013 (c) (1) of DFARs. Notwithstanding the foregoing, the authors grant the U.S. Government and others acting in its behalf permission to use and distribute the software in accordance with the terms specified in this license.

H.18.2. Tix

Copyright (c) 1993-95 Ioi Kim Lam.
Copyright (c) 1996 Expert Interface Technologies.

Permission is hereby granted, without written agreement and without license or royalty fees, to use, copy, modify, and distribute this software and its documentation for any purpose, provided that the above copyright notice and the following two paragraphs appear in all copies of this software.

IN NO EVENT SHALL THE AUTHORS OF THIS SOFTWARE BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE AND ITS DOCUMENTATION, EVEN IF THE AUTHORS OF THIS SOFTWARE HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

THE AUTHORS OF THIS SOFTWARE SPECIFICALLY DISCLAIM ANY WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE SOFTWARE PROVIDED HEREUNDER IS ON AN "AS IS" BASIS, AND THE AUTHORS OF THIS SOFTWARE HAVE NO OBLIGATION TO PROVIDE MAINTENANCE, SUPPORT, UPDATES, ENHANCEMENTS, OR MODIFICATIONS.

Parts of this software are based on the Tcl/Tk software copyrighted by the Regents of the University of California, Sun Microsystems, Inc., and other parties. The original license terms of the Tcl/Tk software distribution is included in the file docs/license.tcltk.

Parts of this software are based on the HTML Library software copyrighted by Sun Microsystems, Inc. The original license terms of the HTML Library software distribution is included in the file docs/license.html_lib.

H.19. pexpect

Python Software Foundation License v2

H.20. phpMyAdmin

Python Software Foundation License v2

H.21. POW

Copyright (c) 2001, 2002, Peter Shannon
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- * Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- * Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- * The name of the contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

H.22. pygtk

GNU Lesser General Public License v2.1

H.23. python

Python Software Foundation License v2

H.24. rcs

GNU General Public License v2

H.25. readline

GNU General Public License v2

H.26. tidy

Copyright (c) 1998-2007 World Wide Web Consortium
(Massachusetts Institute of Technology, European Research
Consortium for Informatics and Mathematics, Keio University).
All Rights Reserved.

This software and documentation is provided "as is," and the copyright holders and contributing author(s) make no representations or warranties, express or implied, including but not limited to, warranties of merchantability or fitness for any particular purpose or that the use of the software or documentation will not infringe any third party patents, copyrights, trademarks or other rights.

The copyright holders and contributing author(s) will not be held liable for any direct, indirect, special or consequential damages arising out of any use of the software or documentation, even if advised of the possibility of such damage.

Permission is hereby granted to use, copy, modify, and distribute this source code, or portions hereof, documentation and executables, for any purpose, without fee, subject to the following restrictions:

1. The origin of this source code must not be misrepresented.
2. Altered versions must be plainly marked as such and must not be misrepresented as being the original source.
3. This Copyright notice may not be removed or altered from any source or altered source distribution.

The copyright holders and contributing author(s) specifically permit, without fee, and encourage the use of this source code as a component for supporting the Hypertext Markup Language in commercial products. If you use this source code in a product, acknowledgment is not required but would be appreciated.

H.27. wget

GNU General Public License v2

Notes

1. <http://git.rockclusters.org>