Chapter 1: Understanding Linux Virtualization

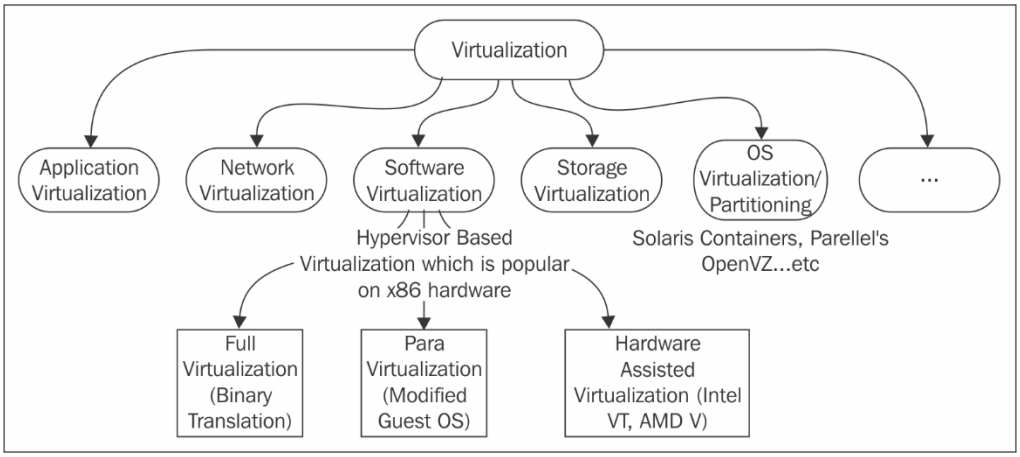

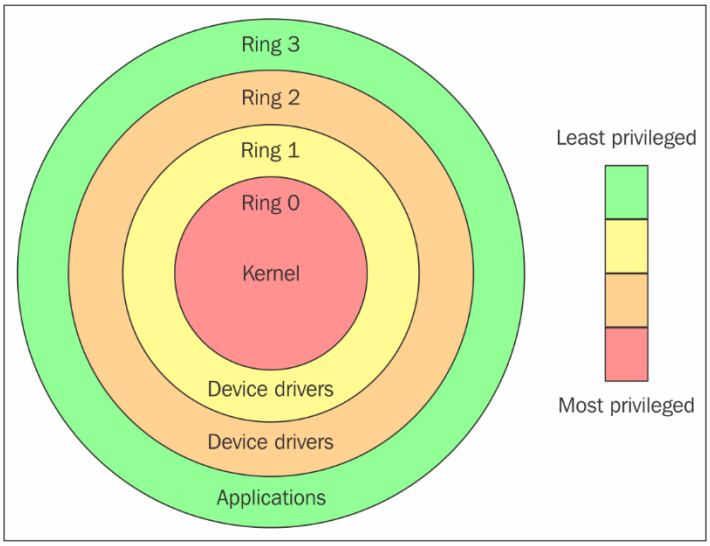

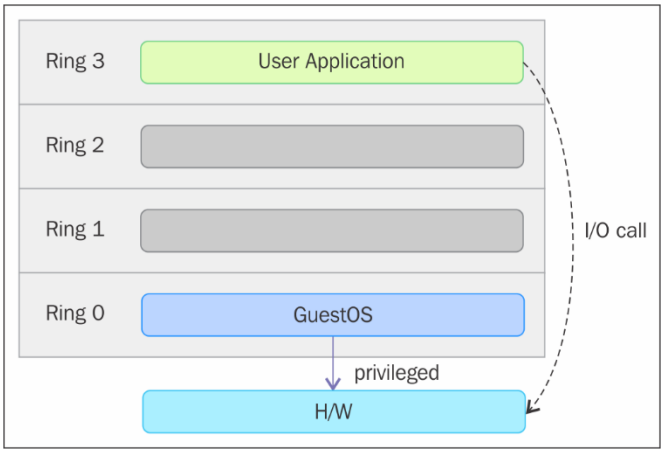

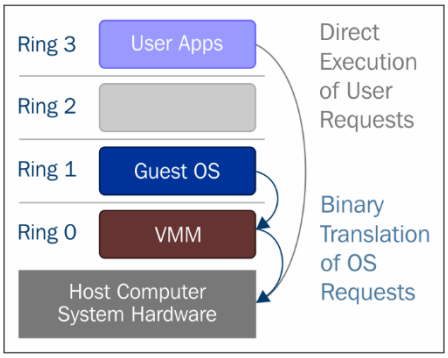
Ring 3
User Apps
Ring 2
Ring 1
Guest OS
Ring 0
VMM
Host Computer
System Hardware
Direct
Execution
of User
Requests
Binary
Translation
of OS
Requests

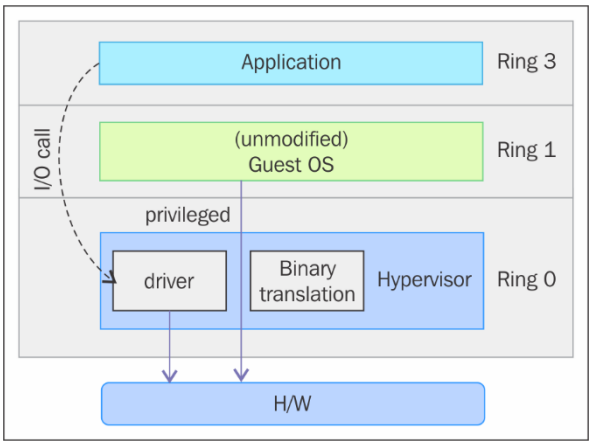
Application
Ring 3
I/O call
(unmodified)
Guest OS
Ring 1
privileged
driver
Binary
translation
Hypervisor
Ring 0
H/W

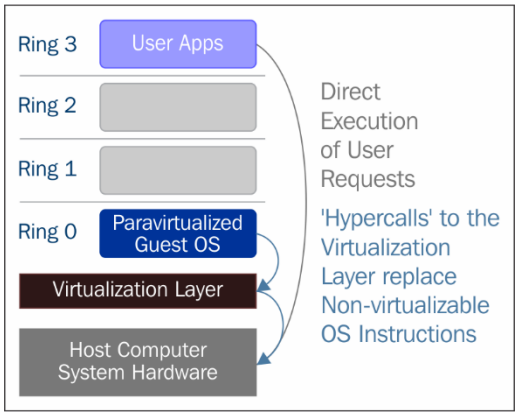
Ring 3
User Apps
Ring 2
Ring 1
Ring 0
Paravirtualized
Guest OS
Virtualization Layer
Host Computer
System Hardware
Direct
Execution
of User
Requests
'Hypercalls' to the
Virtualization
Layer replace
Non-virtualizable
OS Instructions

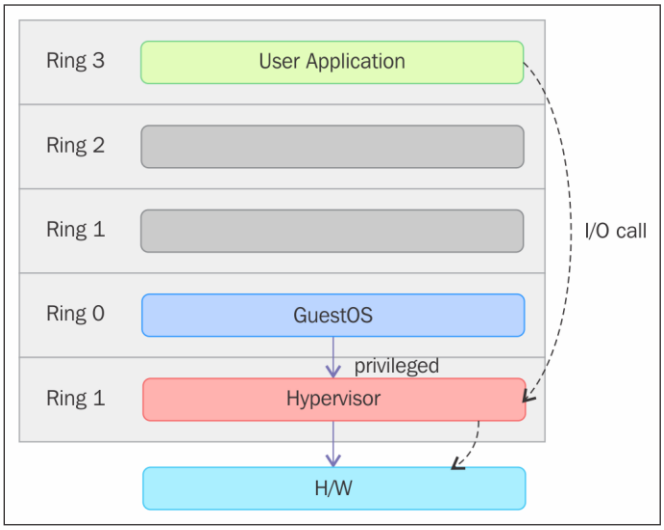
Ring 3
User Application
Ring 2
Ring 1
Ring 0
GuestOS
privileged
Ring 1
Hypervisor
H/W
I/O call

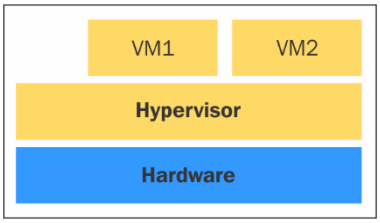
VM1
VM2
Hypervisor
Hardware

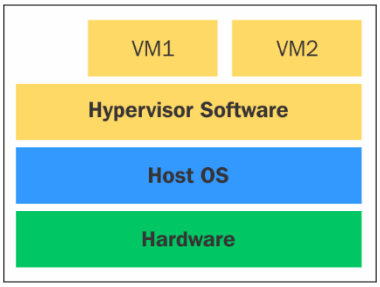
VM1
VM2
Hypervisor Software
Host OS
Hardware

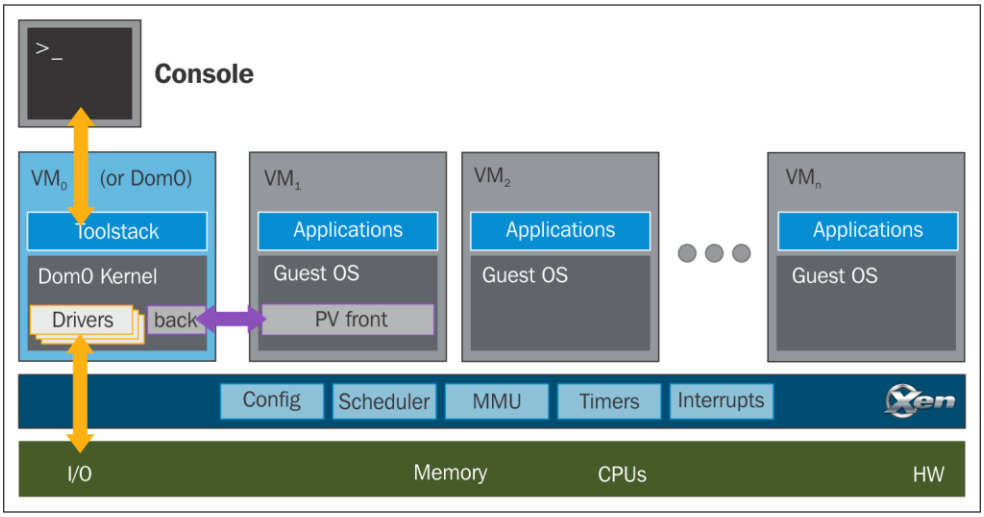
Console
VM0 (or Dom0)
Toolstack
Dom0 Kernel
Drivers
back
VM1
Applications
Guest OS
PV front
VM2
Applications
Guest OS
VMn
Applications
Guest OS
Config
Scheduler
MMU
Timers
Interrupts
Xen
I/O
Memory
CPUs
HW

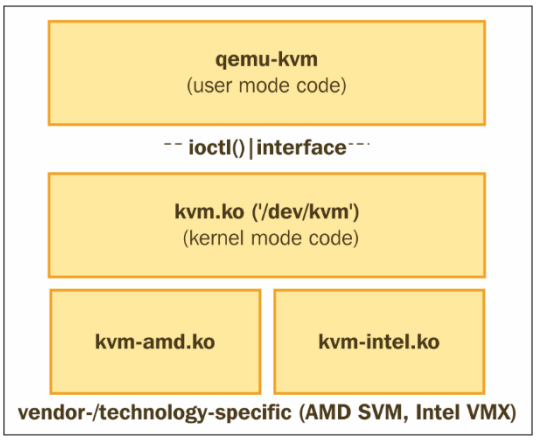
qemu-kvm
(user mode code)
ioctl() interface
kvm.ko ('/dev/kvm')
(kernel mode code)
kvm-amd.ko
kvm-intel.ko
vendor-/technology-specific (AMD SVM, Intel VMX)

Chapter 2: KVM Internals

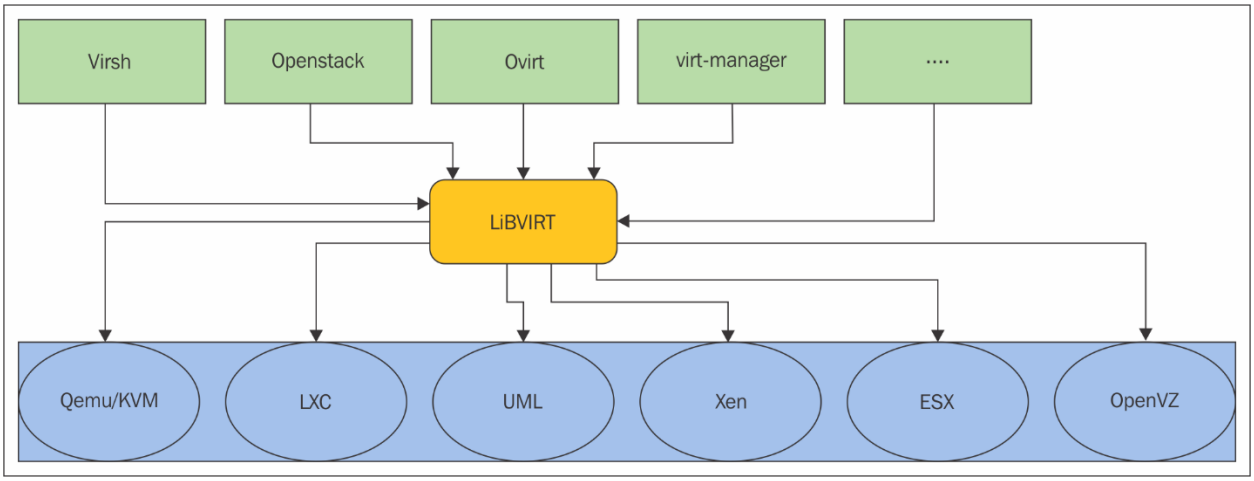

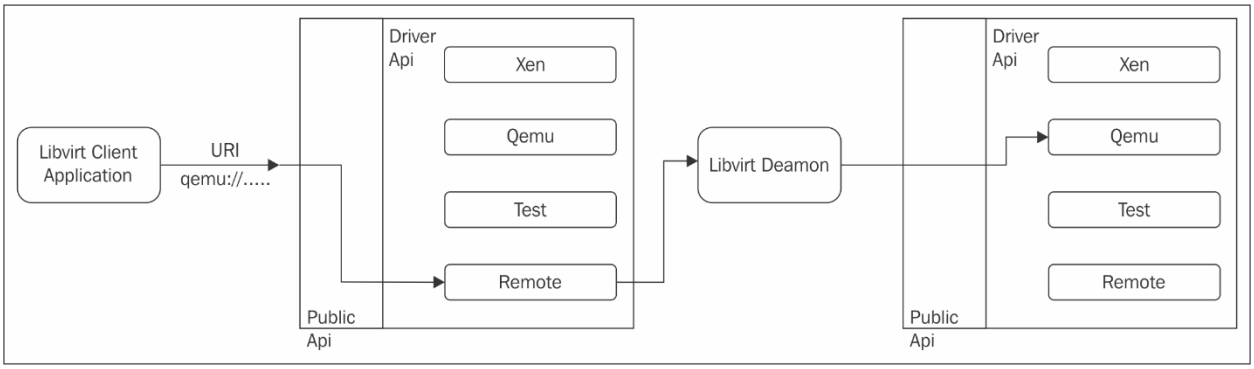


```
[humble-lap ]$ ls
AUTHORS.in       configure.ac     libvirt-admin.pc.in    po
autobuild.sh     COPYING          libvirt-lxc.pc.in      README
autogen.sh       COPYING.LESSER   libvirt.pc.in          README-hacking
bootstrap        daemon           libvirt-qemu.pc.in     run.in
bootstrap.conf   docs             libvirt.spec.in        src
build-aux        examples         m4                     tests
cfg.mk           gnulib           Makefile.am            TODO
ChangeLog-old    HACKING          Makefile.nonreentrant  tools
config-post.h    include          mingw-libvirt.spec.in
[humble-lap ]$
```

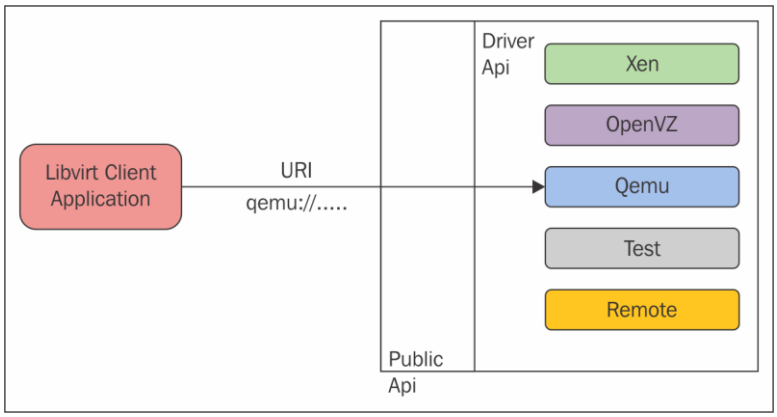

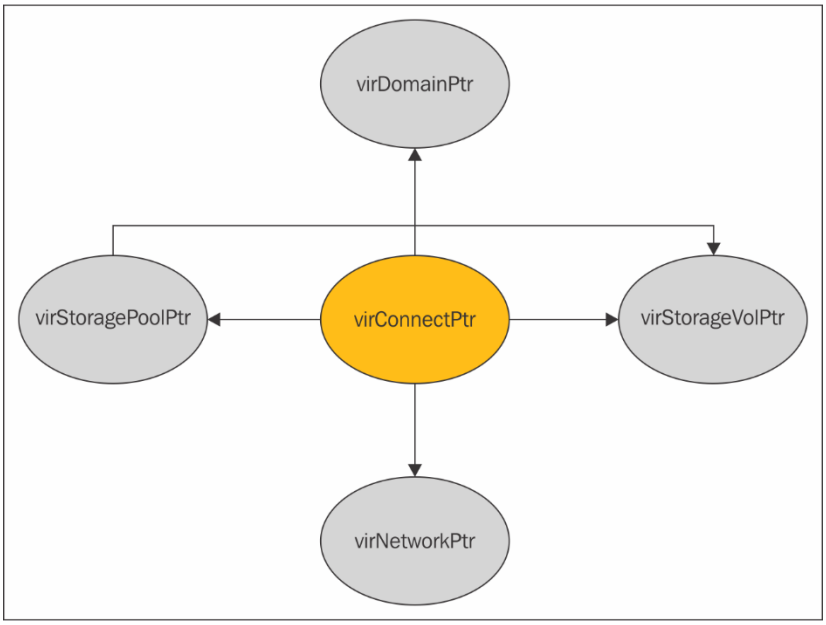


```
[humble-lap ]$ cd src/qemu/
[humble-lap ]$ ls
EVENTHANDLERS.txt    qemu_cgroup.h   qemu_hostdev.h       qemu_monitor_json.h
libvirtd_qemu.aug    qemu_command.c  qemu_hotplug.c       qemu_monitor_text.c
MIGRATION.txt        qemu_command.h  qemu_hotplug.h       qemu_monitor_text.h
qemu_agent.c         qemu.conf       qemu_hotplugpriv.h   qemu_process.c
qemu_agent.h         qemu_conf.c     qemu_interface.c     qemu_process.h
qemu_blockjob.c      qemu_conf.h     qemu_interface.h     qemu_processpriv.h
qemu_blockjob.h      qemu_domain.c   qemu_migration.c     test_libvirtd_qemu.aug.in
qemu_capabilities.c  qemu_domain.h   qemu_migration.h     THREADS.txt
qemu_capabilities.h  qemu_driver.c   qemu_monitor.c
qemu_capspriv.h      qemu_driver.h   qemu_monitor.h
qemu_cgroup.c        qemu_hostdev.c  qemu_monitor_json.c
[humble-lap ]$
```

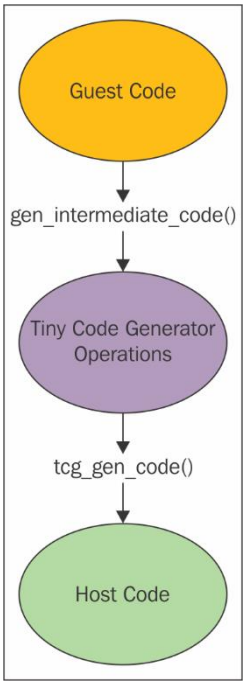
Guest Code
gen_intermediate_code()
Tiny Code Generator Operations
tcg_gen_code()
Host Code

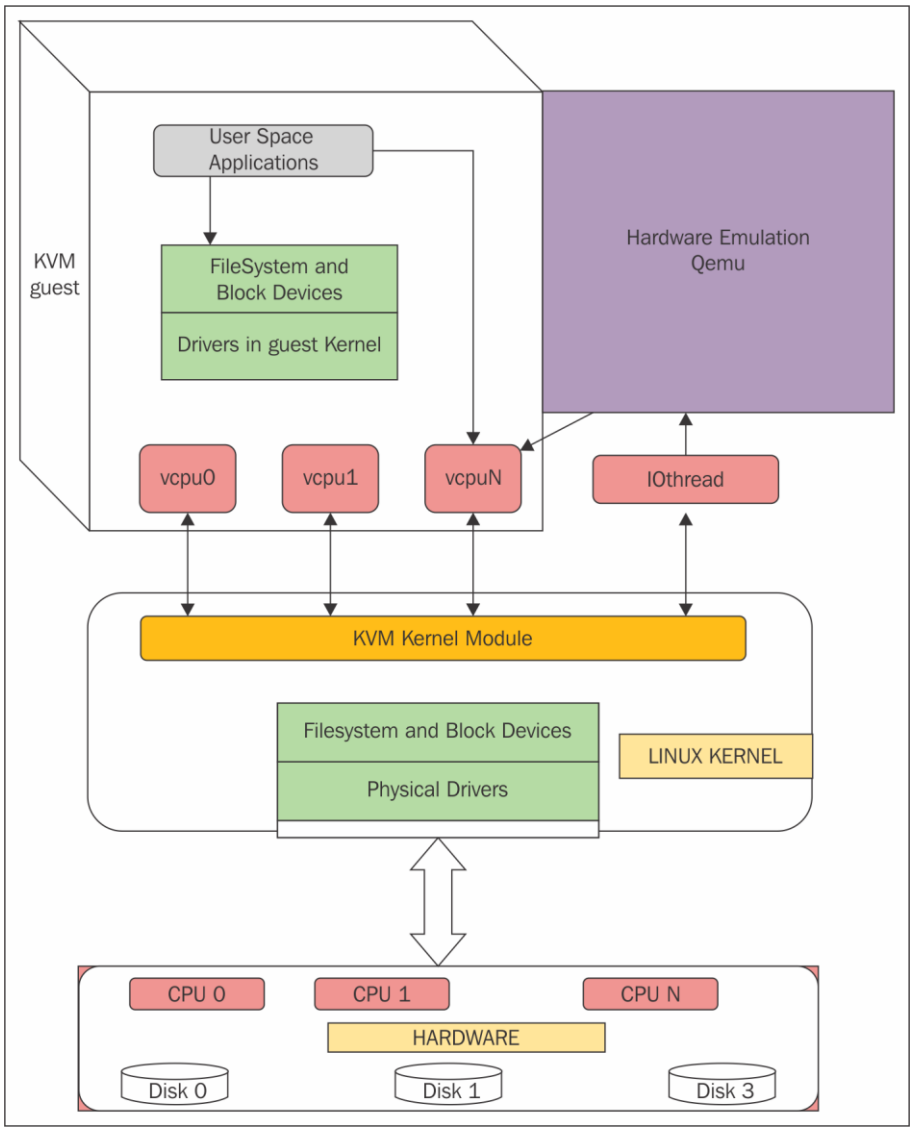
User Space Applications
KVM guest
FileSystem and Block Devices
Drivers in guest Kernel
Hardware Emulation Qemu
vcpu0
vcpu1
vcpuN
IOthread
KVM Kernel Module
Filesystem and Block Devices
Physical Drivers
LINUX KERNEL
CPU 0
CPU 1
CPU N
HARDWARE
Disk 0
Disk 1
Disk 3

```
[ humble-lap ]$ ls
accel.c                 fsdev                   qdev-monitor.c          target-cris
aio-posix.c             gdbstub.c               qdict-test-data.txt     target-i386
aio-win32.c             gdb-xml                 qemu-bridge-helper.c    target-lm32
arch_init.c             HACKING                 qemu-char.c             target-m68k
async.c                 hmp.c                   qemu-doc.texi           target-microblaze
audio                   hmp-commands.hx         qemu-ga.texi            target-mips
backends                hmp-commands-info.hx    qemu-img.c              target-moxie
balloon.c               hmp.h                   qemu-img-cmds.hx        target-openrisc
block                   hw                      qemu-img.texi           target-ppc
block.c                 include                 qemu-io.c               target-s390x
blockdev.c              iohandler.c             qemu-io-cmds.c          target-sh4
blockdev-nbd.c          ioport.c                qemu-log.c              target-sparc
blockjob.c              iothread.c              qemu-nbd.c              target-tilegx
bootdevice.c            kvm-all.c               qemu-nbd.texi           target-tricore
bsd-user                kvm-stub.c              qemu.nsi                target-unicore32
bt-host.c               libdecnumber            qemu-options.h          target-xtensa
bt-vhci.c               LICENSE                 qemu-options.hx         tcg
Changelog               linux-headers           qemu-options-wrapper.h  tcg-runtime.c
CODING_STYLE            linux-user              qemu.sasl               tci.c
configure               main-loop.c             qemu-seccomp.c          tests
contrib                 MAINTAINERS             qemu-tech.texi          thread-pool.c
COPYING                 Makefile                qemu-timer.c            thunk.c
COPYING.LIB             Makefile.objs           qga                     tpm.c
cpu-exec.c              Makefile.target         qjson.c                 trace
cpu-exec-common.c       memory.c                qmp.c                   trace-events
cpus.c                  memory_mapping.c        qmp-commands.hx         translate-all.c
cputlb.c                migration               qobject                 translate-all.h
crypto                  module-common.c         qom                     translate-common.c
cscope.out              monitor.c               qtest.c                 ui
default-configs         nbd.c                   README                  user-exec.c
device-hotplug.c        net                     replay                  util
device_tree.c           numa.c                  roms                    VERSION
disas                   os-posix.c              rules.mak               version.rc
disas.c                 os-win32.c              scripts                 vl.c
dma-helpers.c           page_cache.c            slirp                   xen-common.c
docs                    pc-bios                 softmmu_template.h      xen-common-stub.c
dtc                     pixman                  spice-qemu-char.c       xen-hvm.c
dump.c                  po                      stubs                   xen-hvm-stub.c
exec.c                  qapi                    target-alpha            xen-mapcache.c
fpu                     qapi-schema.json        target-arm
```

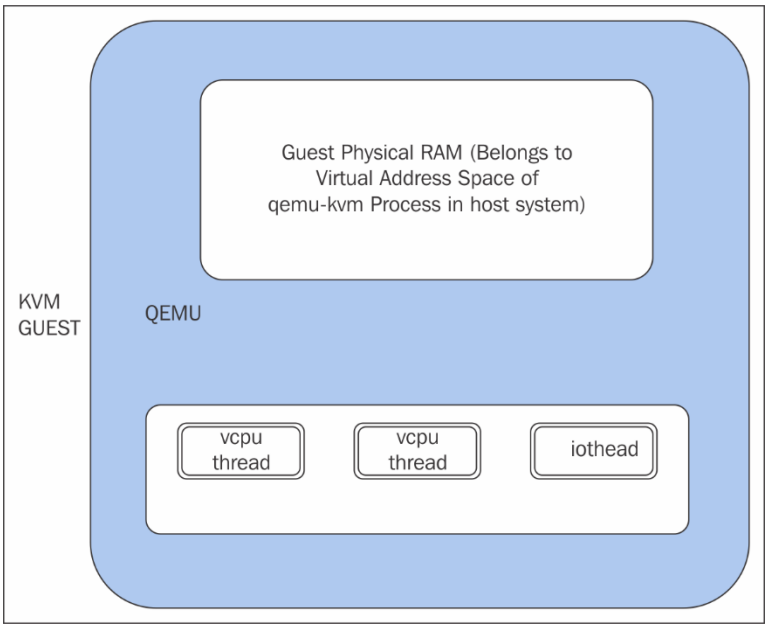
Guest Physical RAM (Belongs to
Virtual Address Space of
qemu-kvm Process in host system)
KVM
GUEST
QEMU
vcpu
thread
vcpu
thread
iothead

Chapter 3: Setting Up Standalone KVM Virtualization

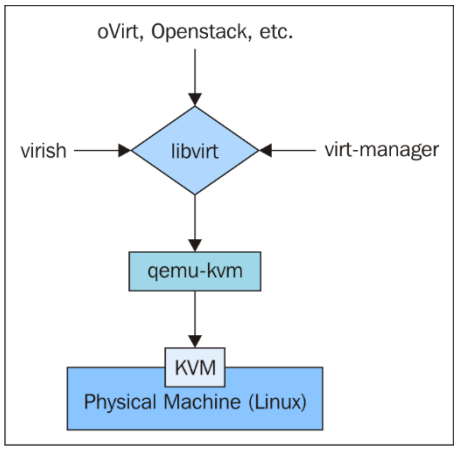


```
[root@kvmHOST ~]# yum repolist
Loaded plugins: langpacks
repo id                          repo name                                status
fedora/21/x86_64                 Fedora 21 - x86_64                       42,816
updates/21/x86_64                Fedora 21 - x86_64 - Updates             16,716
repolist: 59,532
[root@kvmHOST ~]#
```

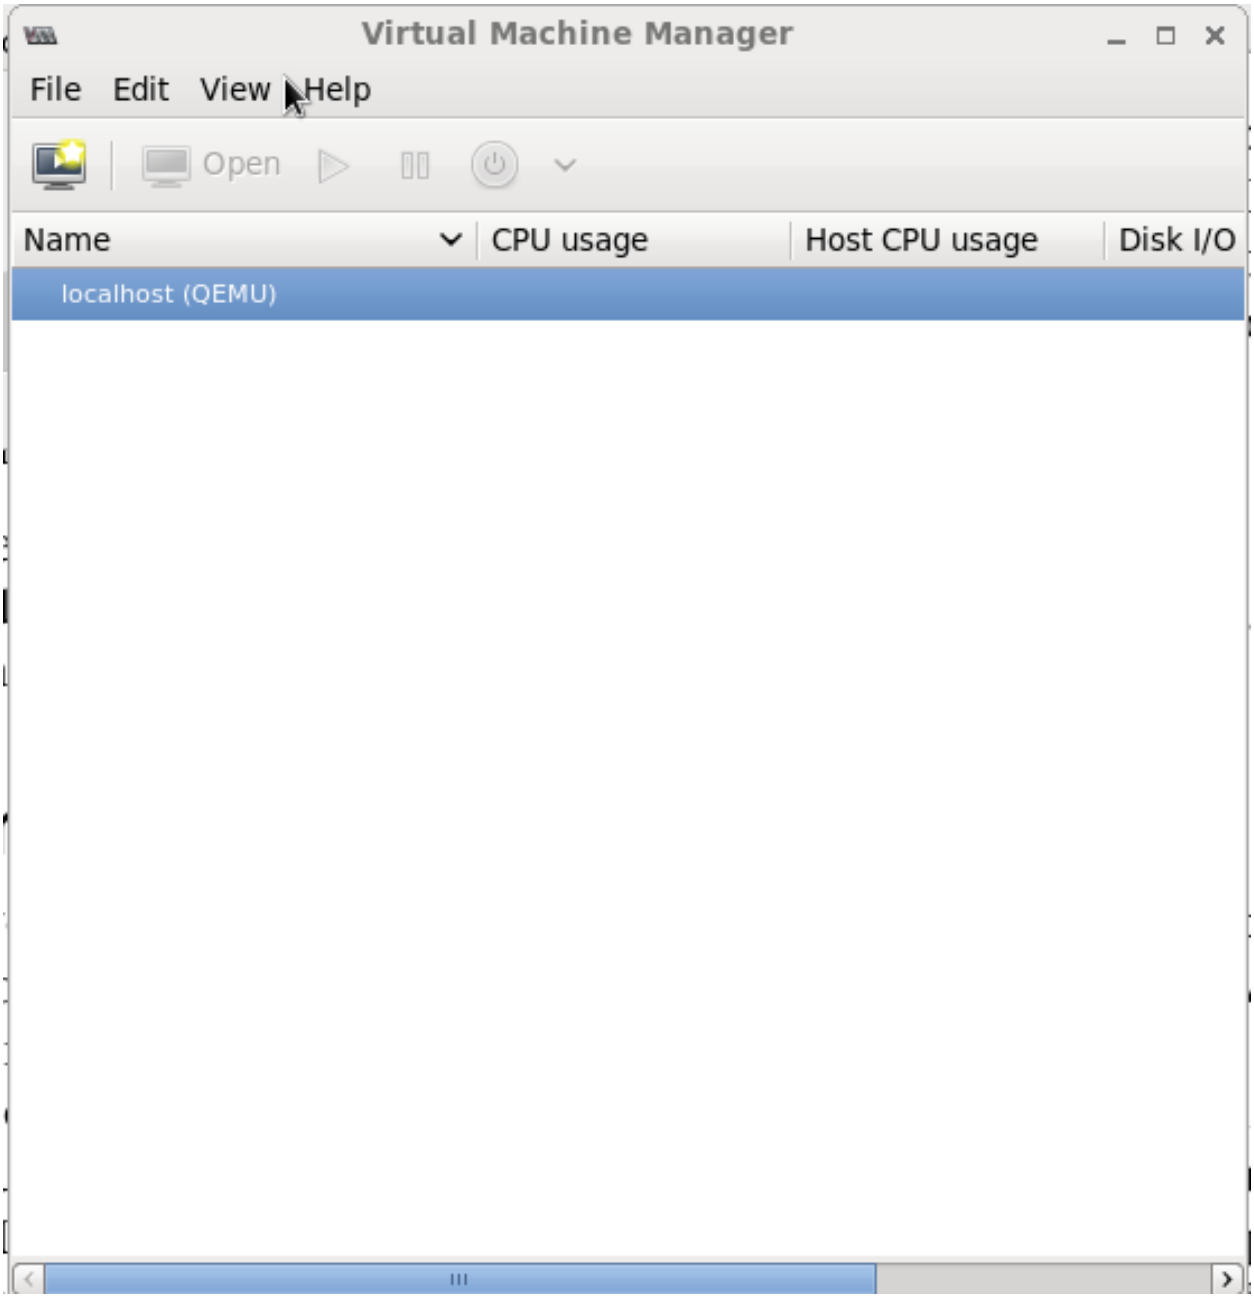
Virtual Machine Manager
File
Edit
View
Help
Open
Name
CPU usage
Host CPU usage
Disk I/O
localhost (QEMU)

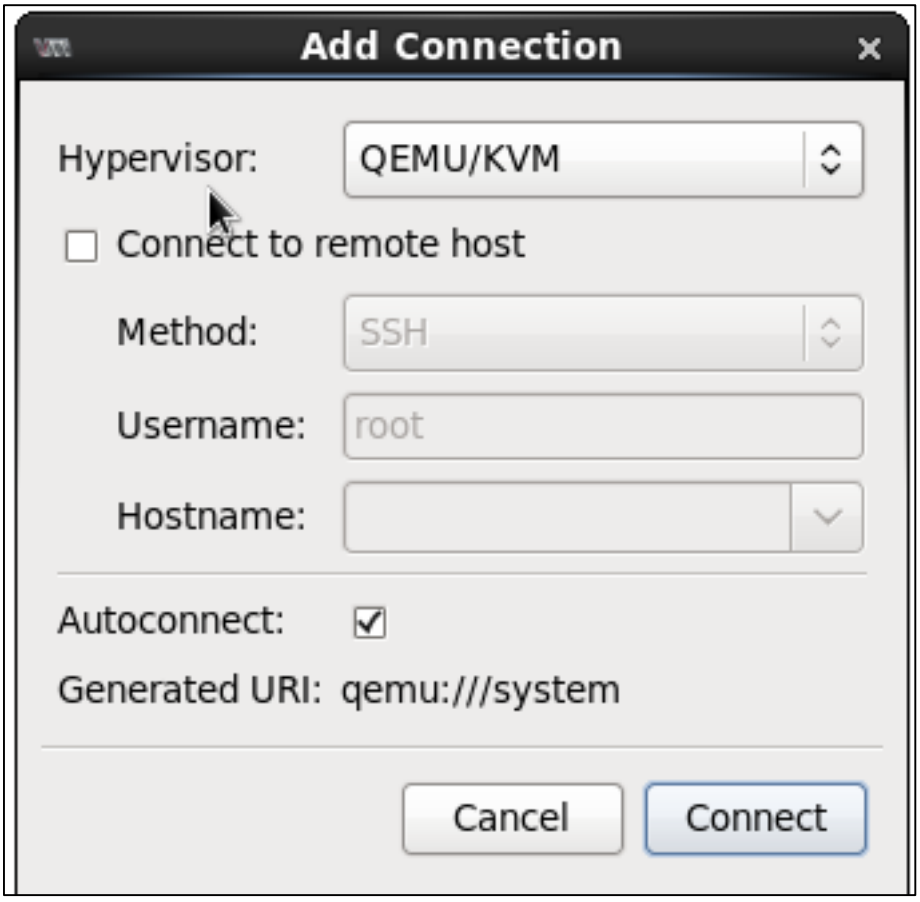
Add Connection
Hypervisor:
QEMU/KVM
Connect to remote host
Method:
SSH
Username:
root
Hostname:
Autoconnect:
Generated URI: qemu:///system
Cancel
Connect

Chapter 4: Getting Started with libvirt and Creating Your First Virtual Machines

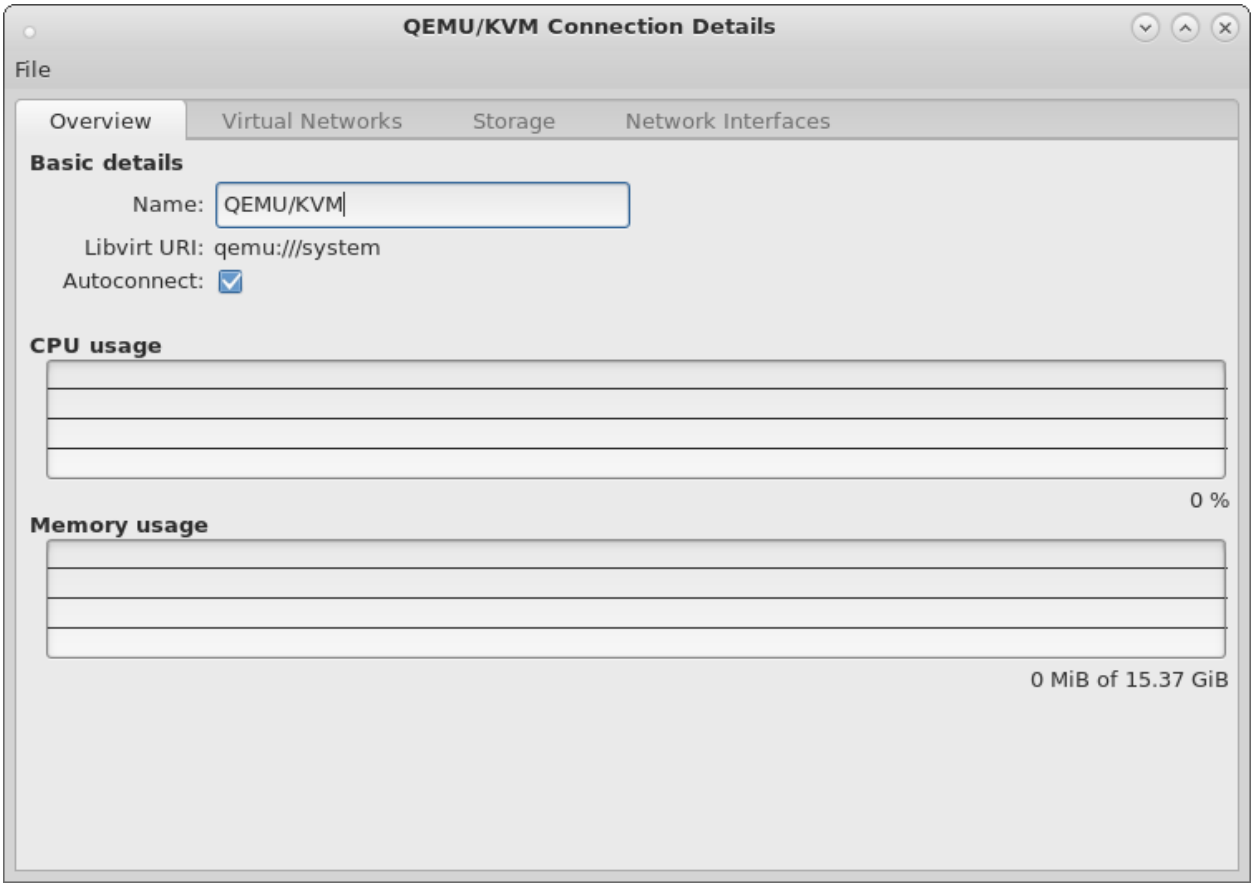

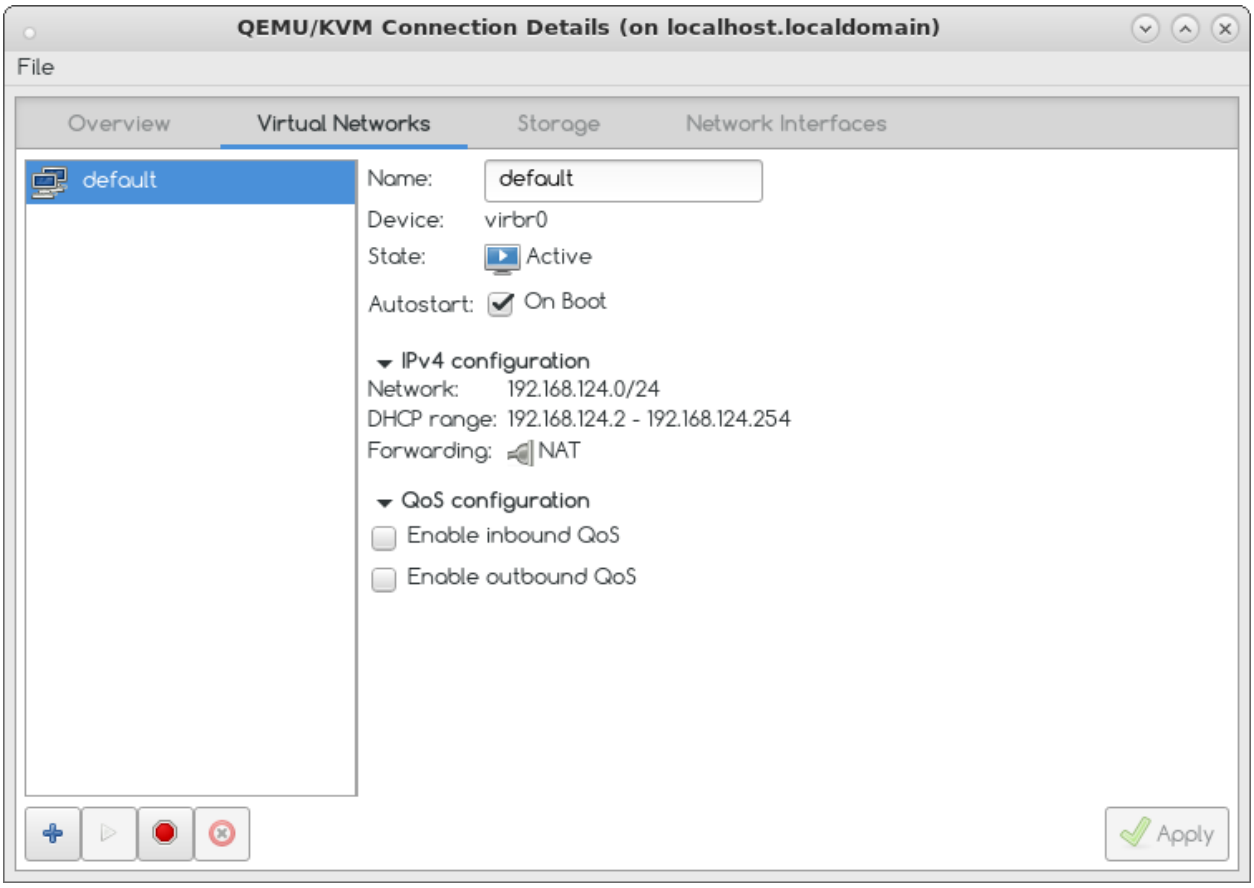
QEMU/KVM Connection Details (on localhost.localdomain)
File
Overview
Virtual Networks
Storage
Network Interfaces
default
Name:
default
Device:
virbr0
State:
Active
Autostart:
On Boot
IPv4 configuration
Network:
192.168.124.0/24
DHCP range:
192.168.124.2 - 192.168.124.254
Forwarding:
NAT
QoS configuration
Enable inbound QoS
Enable outbound QoS
Apply

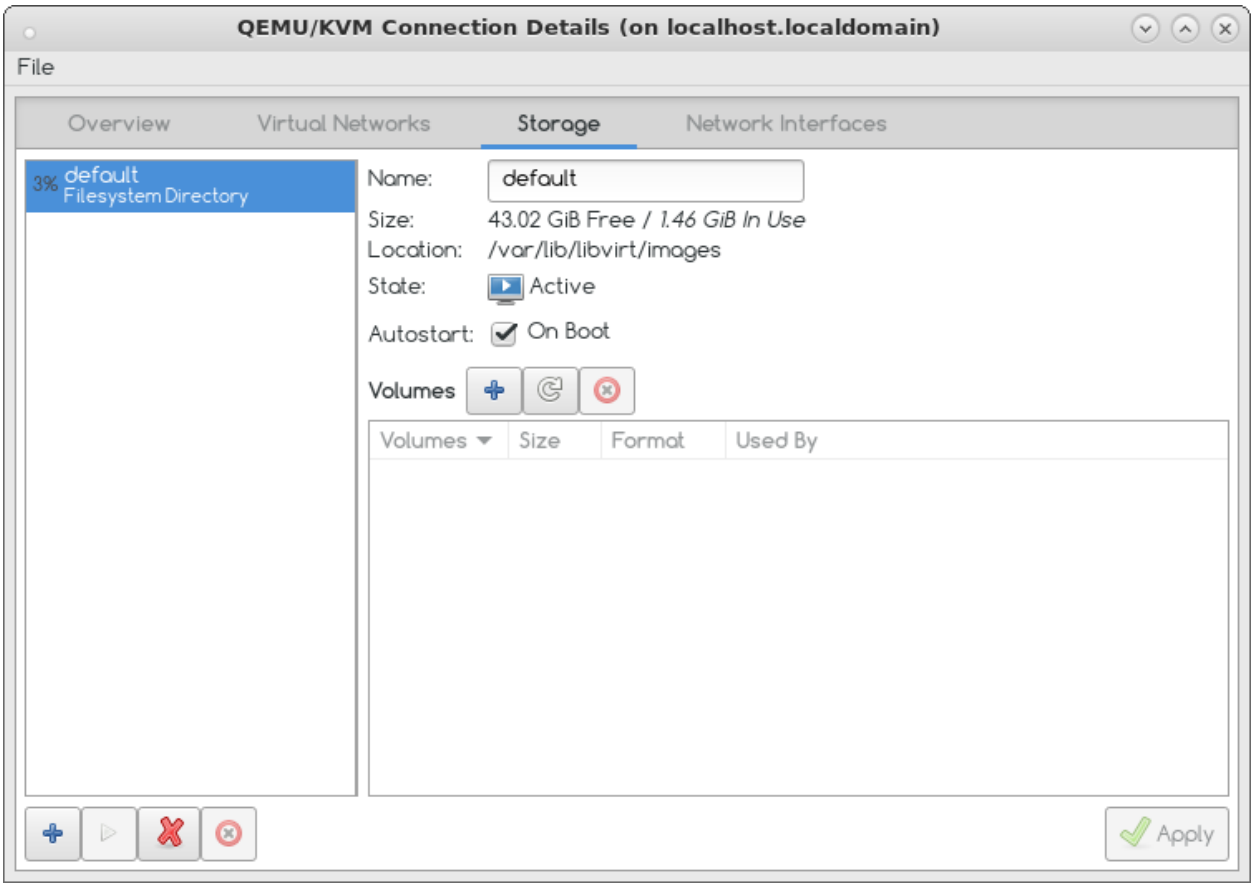
QEMU/KVM Connection Details (on localhost.localdomain)
File
Overview
Virtual Networks
Storage
Network Interfaces
3% default
Filesystem Directory
Name:
default
Size:
43.02 GiB Free / 1.46 GiB In Use
Location:
/var/lib/libvirt/images
State:
Active
Autostart:
On Boot
Volumes
Volumes
Size
Format
Used By
Apply

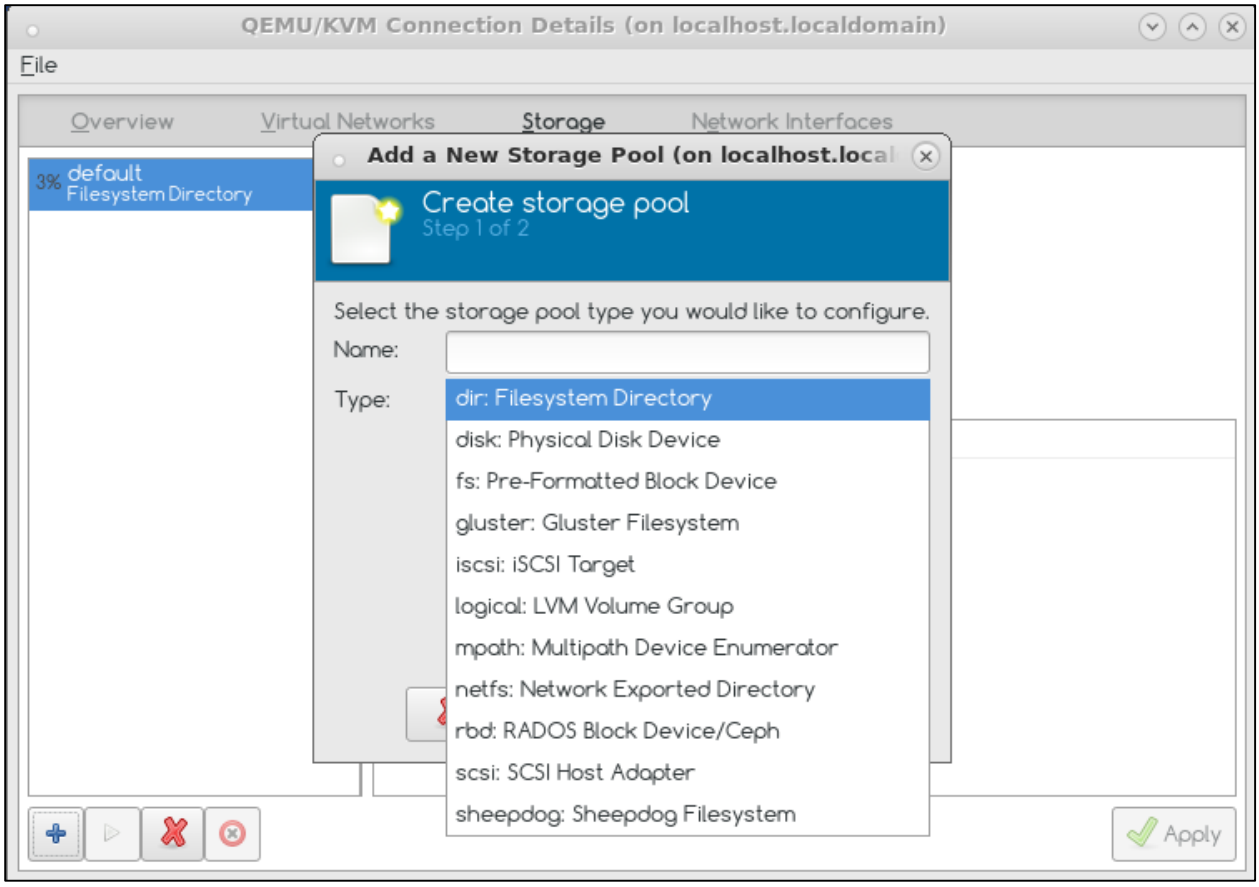
QEMU/KVM Connection Details (on localhost.localdomain)
File
Overview
Virtual Networks
Storage
Network Interfaces
3% default
Filesystem Directory
Add a New Storage Pool (on localhost.local
Create storage pool
Step 1 of 2
Select the storage pool type you would like to configure.
Name:
Type:
dir: Filesystem Directory
disk: Physical Disk Device
fs: Pre-Formatted Block Device
gluster: Gluster Filesystem
iscsi: iSCSI Target
logical: LVM Volume Group
mpath: Multipath Device Enumerator
netfs: Network Exported Directory
rbd: RADOS Block Device/Ceph
scsi: SCSI Host Adapter
sheepdog: Sheepdog Filesystem
Apply

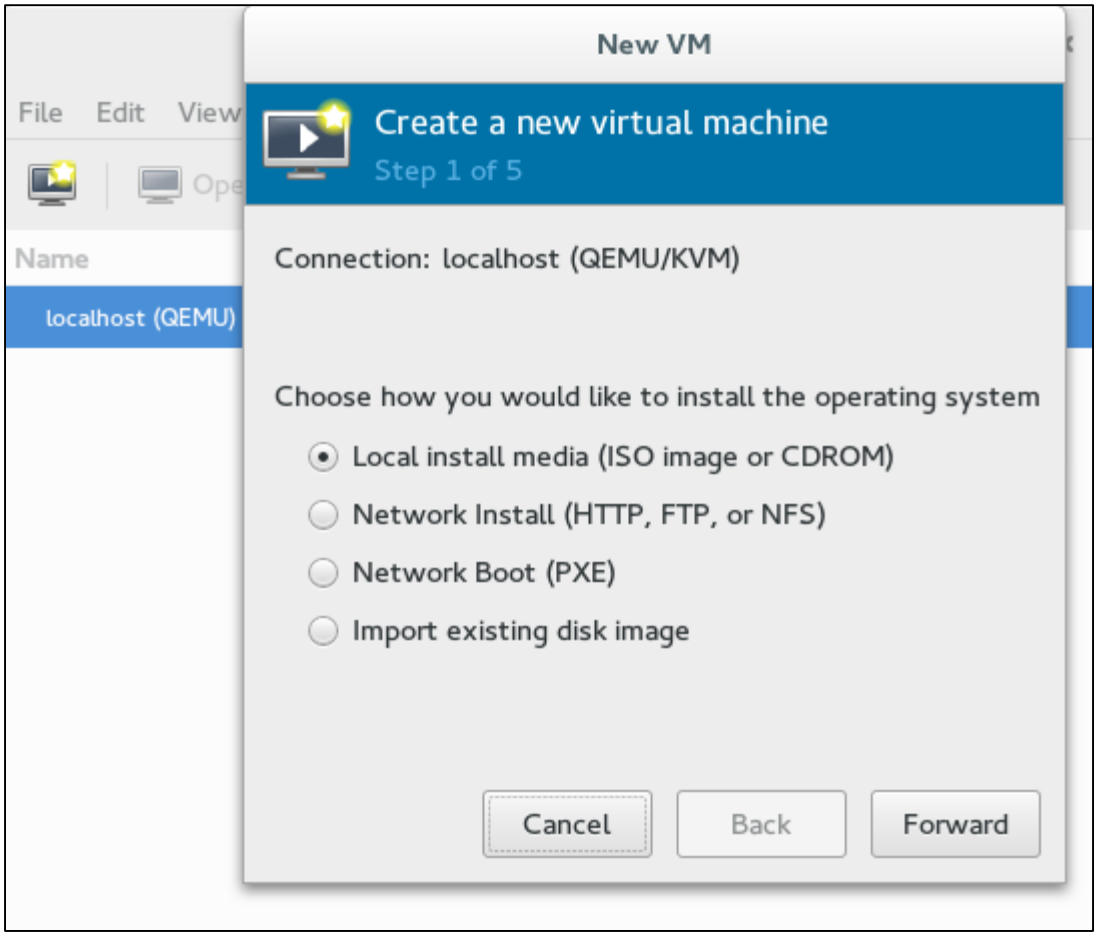
New VM
File
Edit
View
Name
localhost (QEMU)
Create a new virtual machine
Step 1 of 5
Connection: localhost (QEMU/KVM)
Choose how you would like to install the operating system
Local install media (ISO image or CDROM)
Network Install (HTTP, FTP, or NFS)
Network Boot (PXE)
Import existing disk image
Cancel
Back
Forward

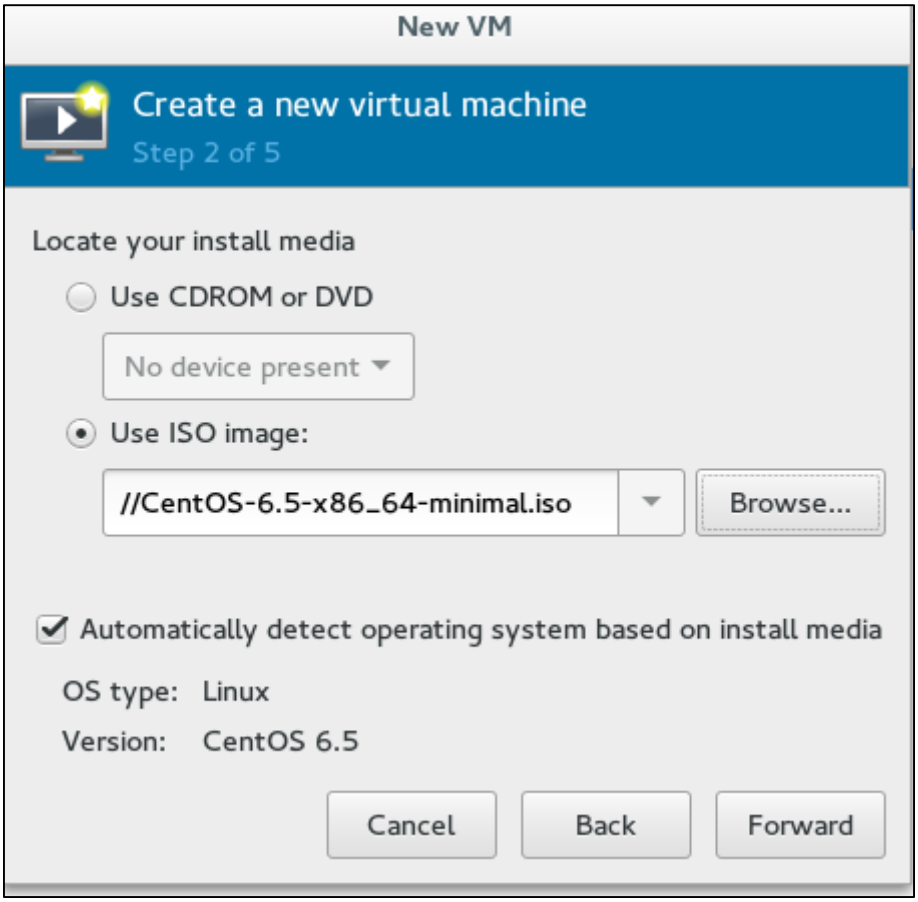
New VM
Create a new virtual machine
Step 2 of 5
Locate your install media
Use CDROM or DVD
No device present
Use ISO image:
//CentOS-6.5-x86_64-minimal.iso
Browse...
Automatically detect operating system based on install media
OS type: Linux
Version: CentOS 6.5
Cancel
Back
Forward

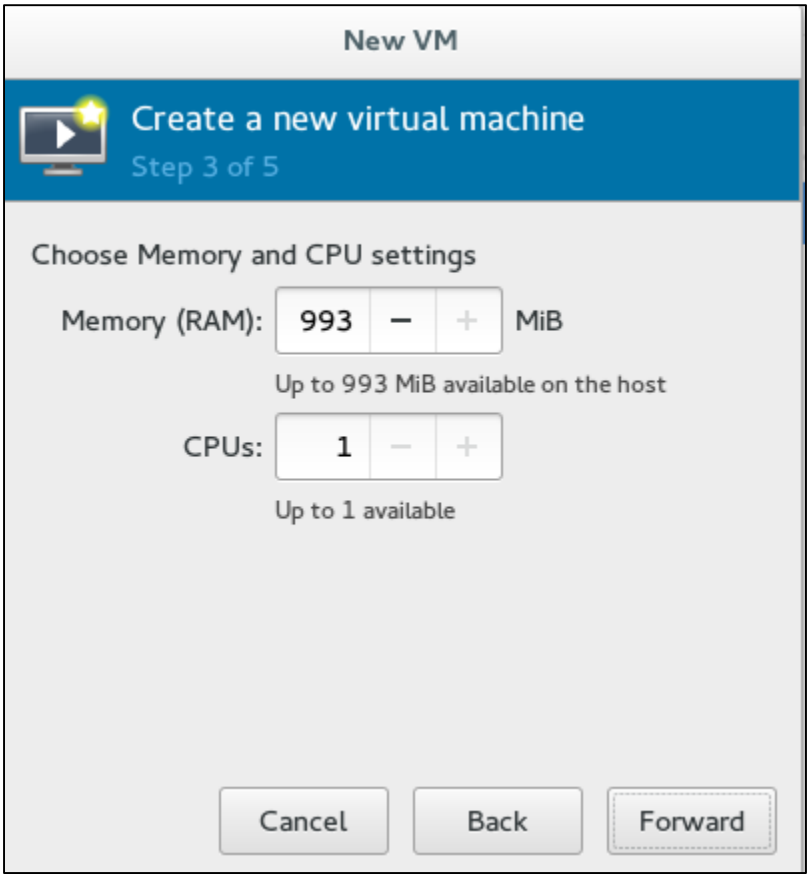
New VM
Create a new virtual machine
Step 3 of 5
Choose Memory and CPU settings
Memory (RAM):
993
MiB
Up to 993 MiB available on the host
CPUs:
1
Up to 1 available
Cancel
Back
Forward

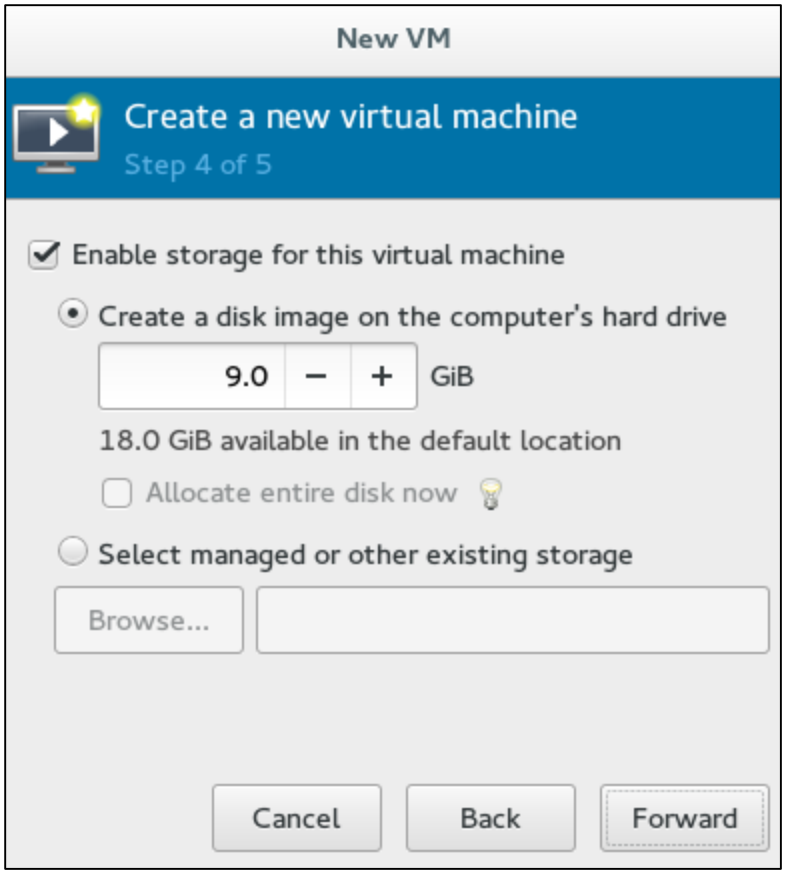
New VM
Create a new virtual machine
Step 4 of 5
Enable storage for this virtual machine
Create a disk image on the computer's hard drive
9.0
GiB
18.0 GiB available in the default location
Allocate entire disk now
Select managed or other existing storage
Browse...
Cancel
Back
Forward

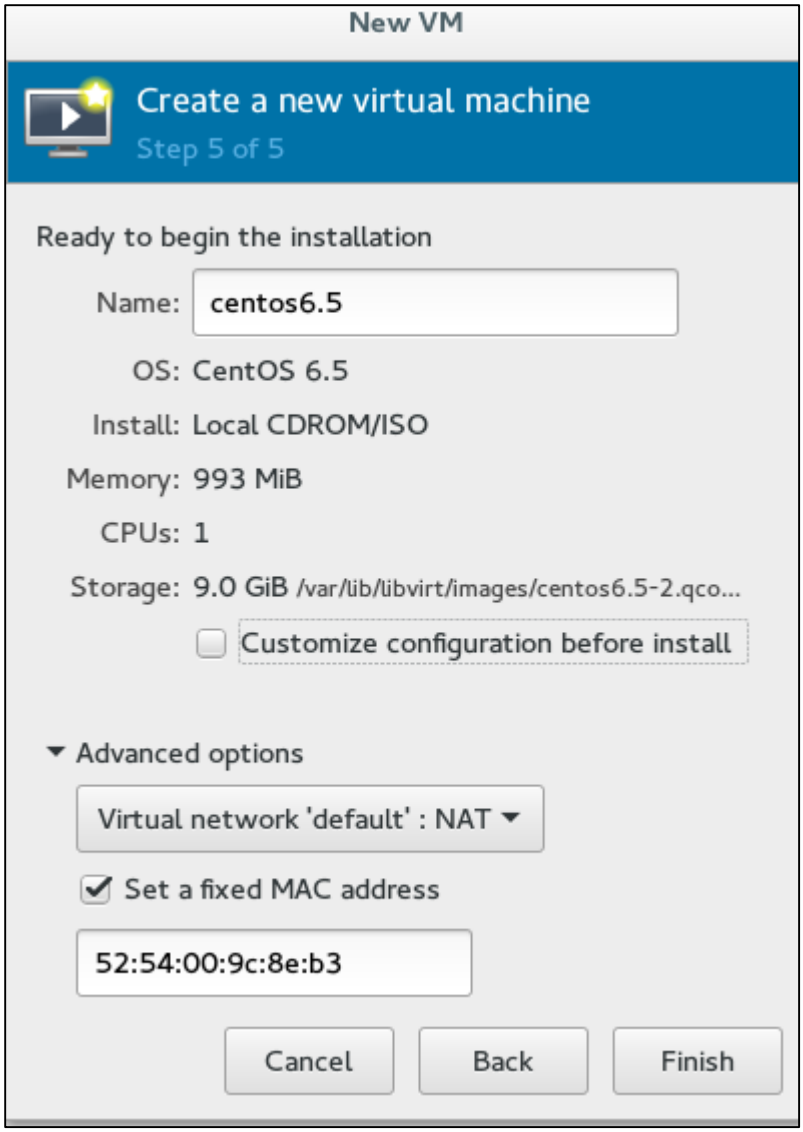
New VM
Create a new virtual machine
Step 5 of 5
Ready to begin the installation
Name: centos6.5
OS: CentOS 6.5
Install: Local CDROM/ISO
Memory: 993 MiB
CPUs: 1
Storage: 9.0 GiB /var/lib/libvirt/images/centos6.5-2.qco...
Customize configuration before install
Advanced options
Virtual network 'default' : NAT
Set a fixed MAC address
52:54:00:9c:8e:b3
Cancel
Back
Finish

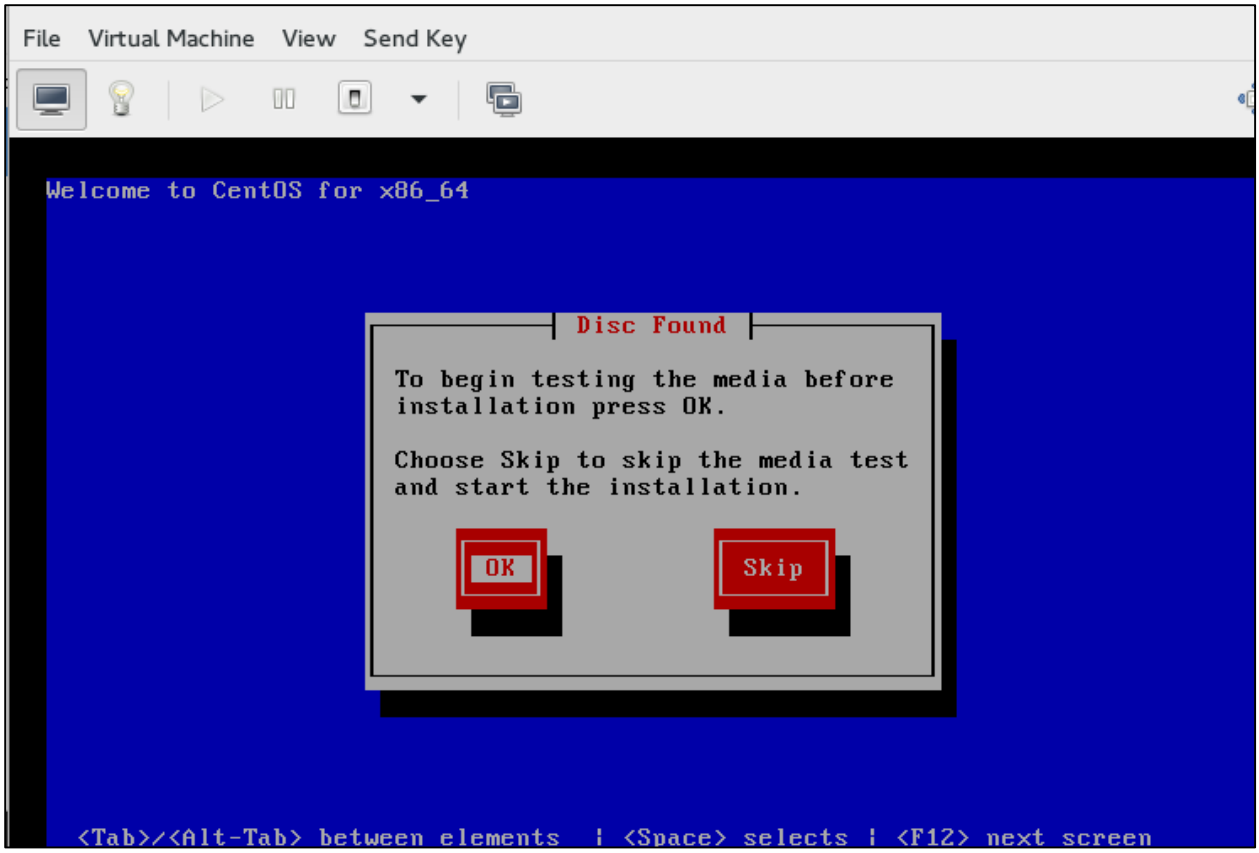
File
Virtual Machine
View
Send Key
Welcome to CentOS for x86_64
Disc Found
To begin testing the media before installation press OK.
Choose Skip to skip the media test and start the installation.
OK
Skip
<Tab>/<Alt-Tab> between elements | <Space> selects | <F12> next screen

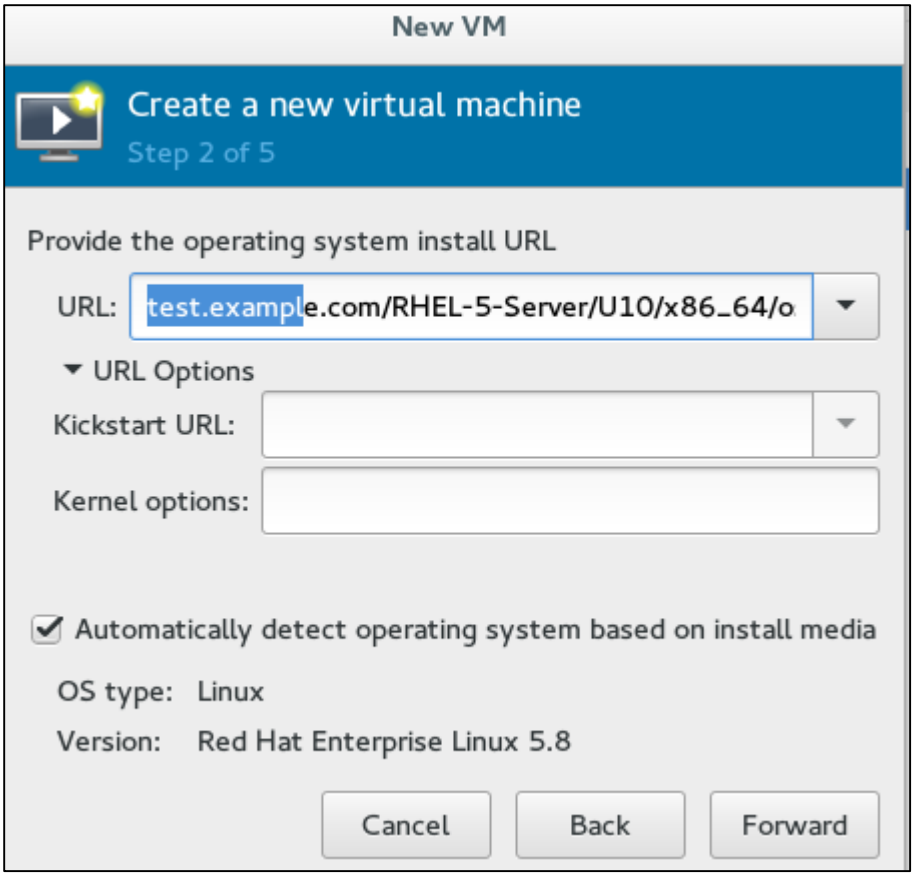
New VM
Create a new virtual machine
Step 2 of 5
Provide the operating system install URL
URL:
test.example.com/RHEL-5-Server/U10/x86_64/o
URL Options
Kickstart URL:
Kernel options:
Automatically detect operating system based on install media
OS type: Linux
Version: Red Hat Enterprise Linux 5.8
Cancel
Back
Forward

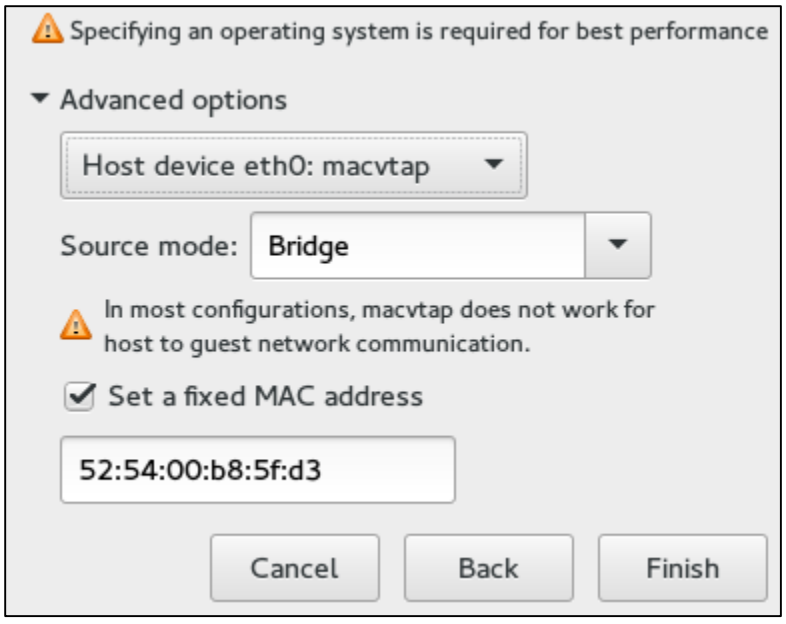
Specifying an operating system is required for best performance
Advanced options
Host device eth0: macvtap
Source mode:
Bridge
In most configurations, macvtap does not work for host to guest network communication.
Set a fixed MAC address
52:54:00:b8:5f:d3
Cancel
Back
Finish

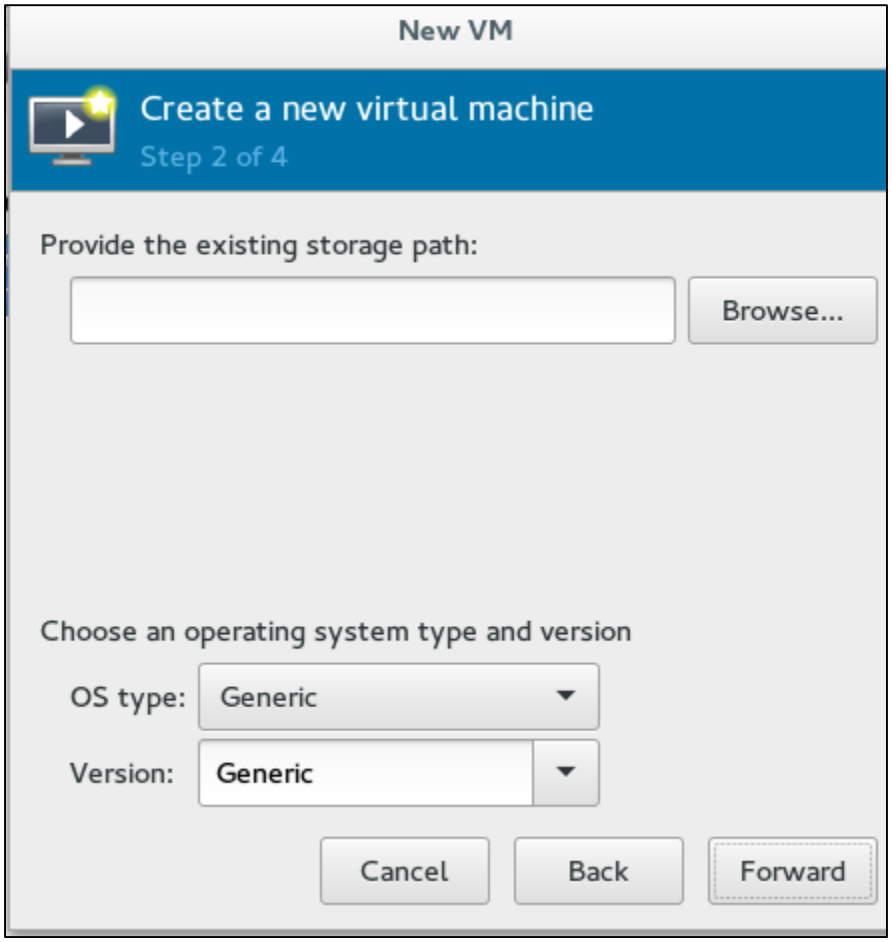
New VM
Create a new virtual machine
Step 2 of 4
Provide the existing storage path:
Browse...
Choose an operating system type and version
OS type:
Generic
Version:
Generic
Cancel
Back
Forward

Chapter 5: Network and Storage

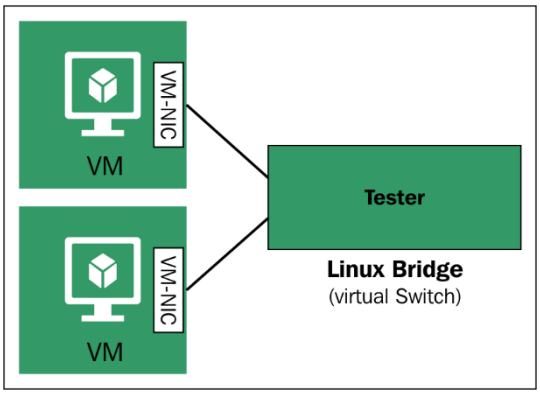

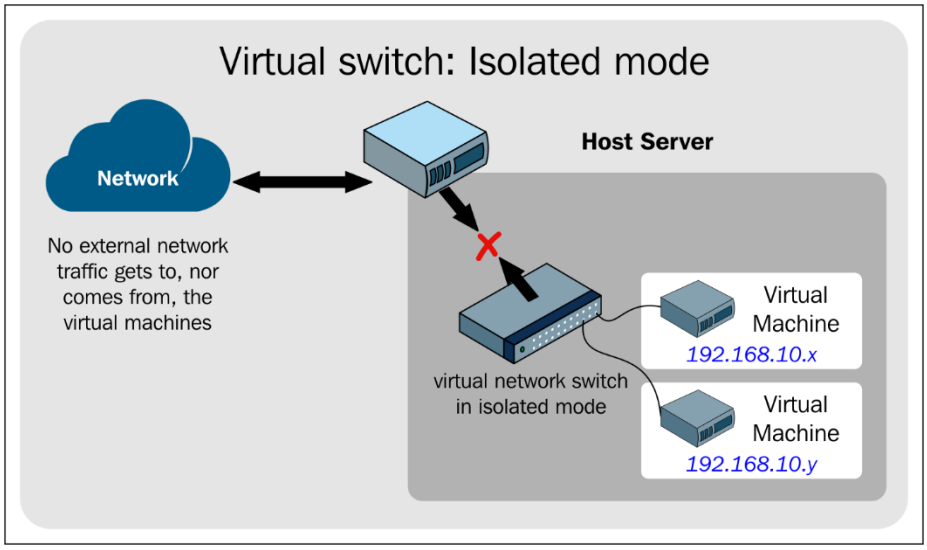

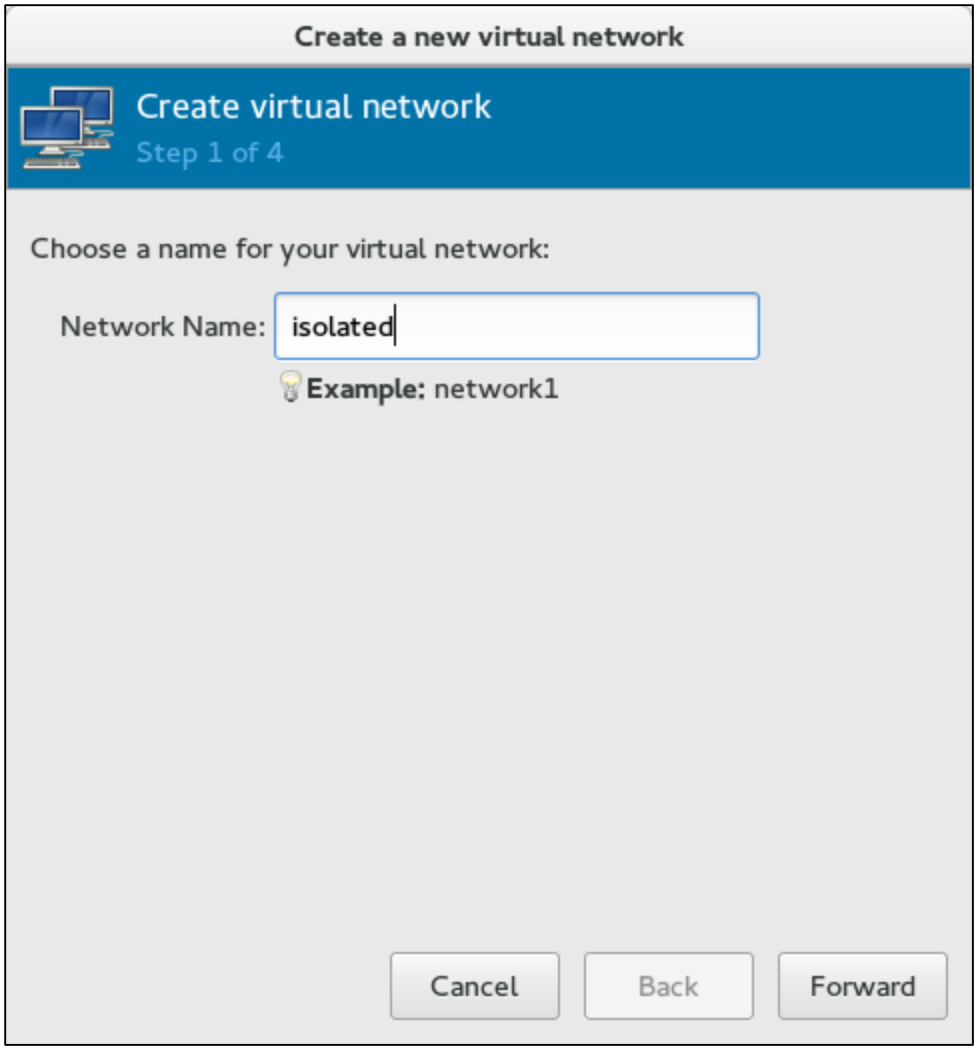
Create a new virtual network
Create virtual network
Step 1 of 4
Choose a name for your virtual network:
Network Name:
isolated
Example: network1
Cancel
Back
Forward

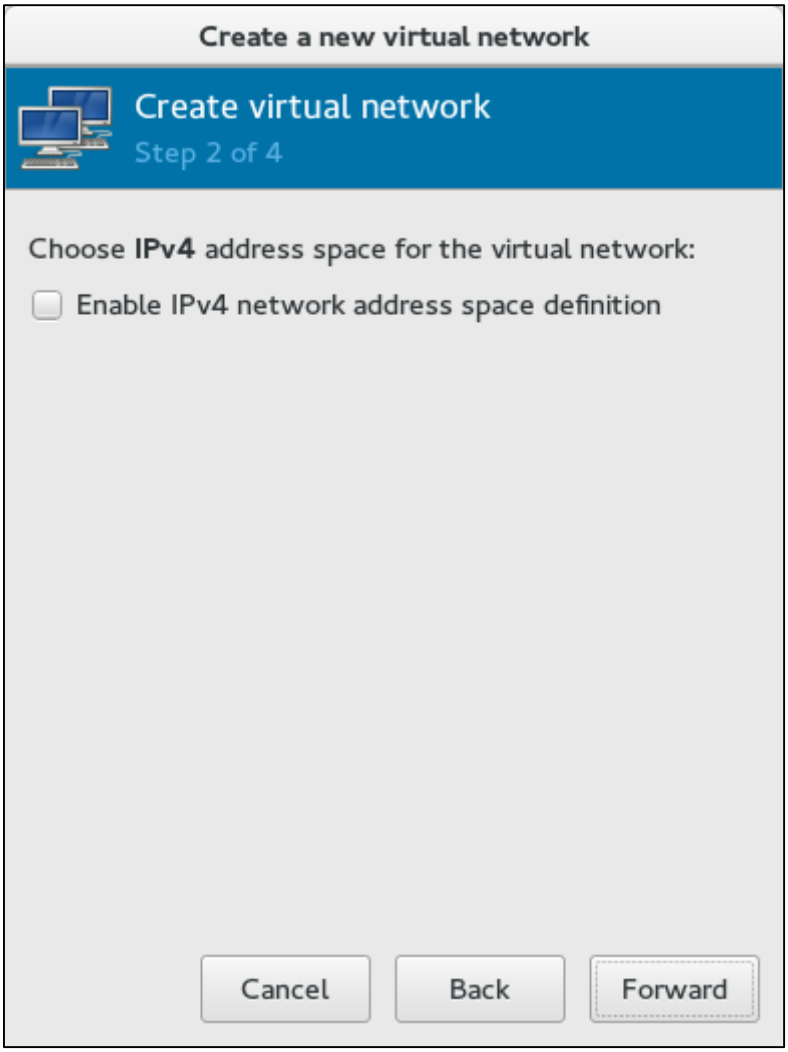
Create a new virtual network
Create virtual network
Step 2 of 4
Choose IPv4 address space for the virtual network:
Enable IPv4 network address space definition
Cancel
Back
Forward

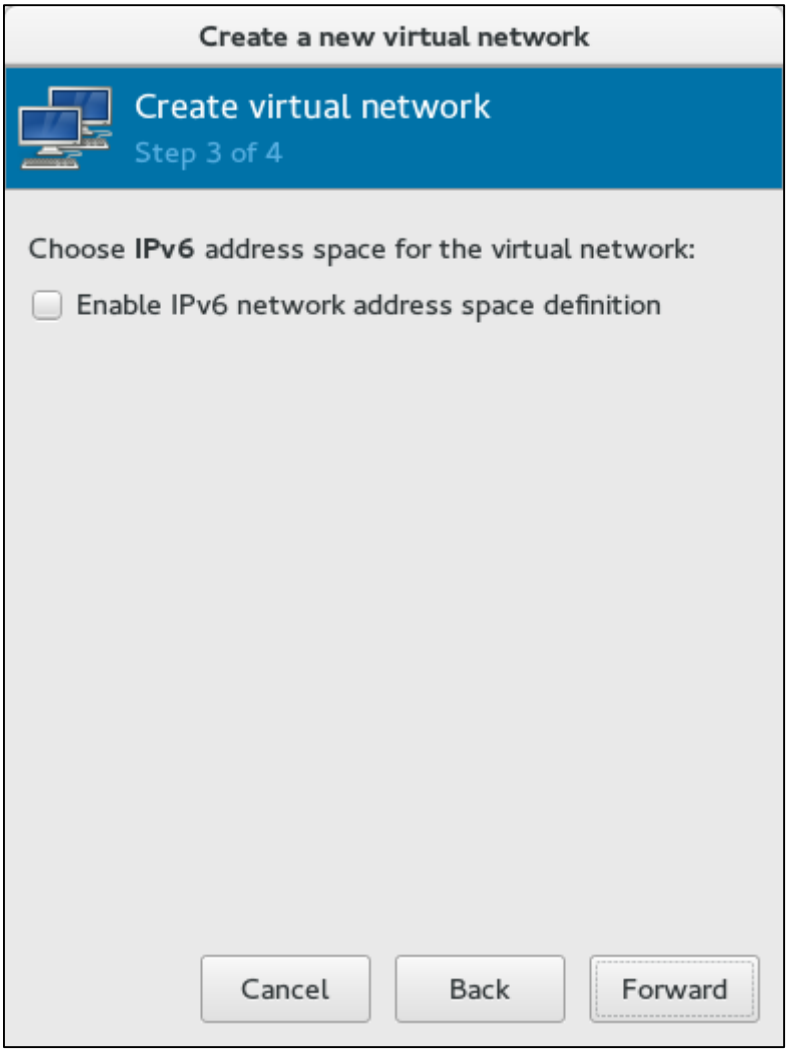
Create a new virtual network
Create virtual network
Step 3 of 4
Choose IPv6 address space for the virtual network:
Enable IPv6 network address space definition
Cancel
Back
Forward

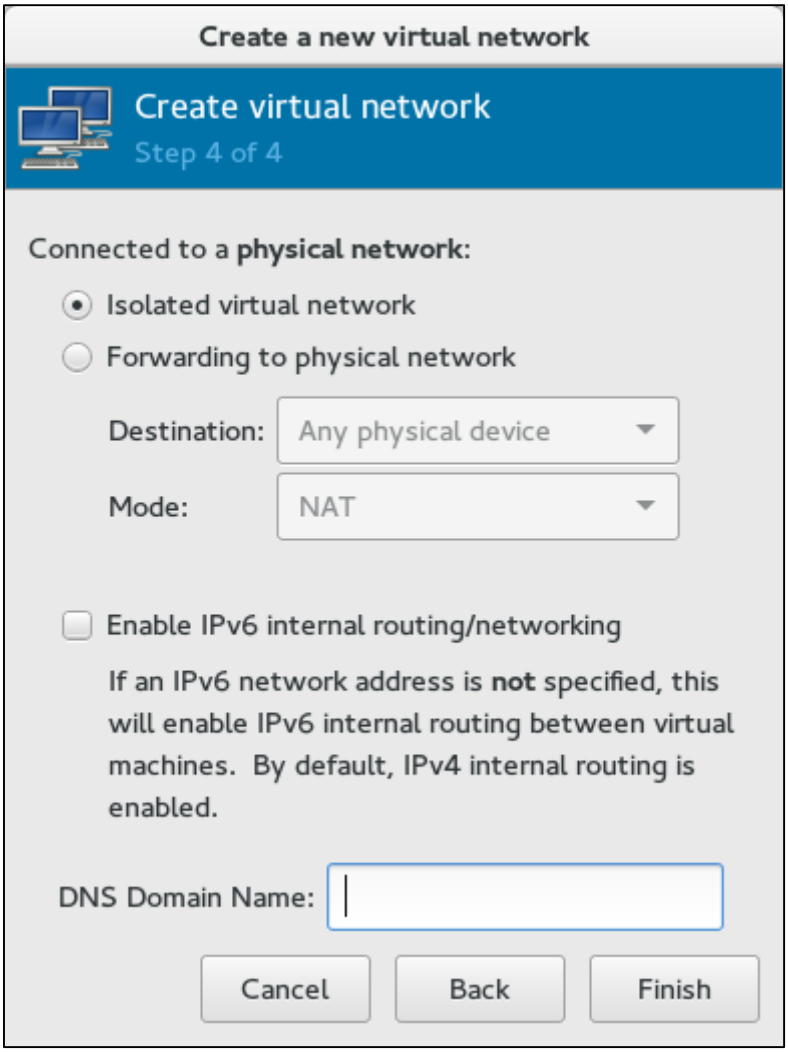
Create a new virtual network
Create virtual network
Step 4 of 4
Connected to a physical network:
Isolated virtual network
Forwarding to physical network
Destination:
Any physical device
Mode:
NAT
Enable IPv6 internal routing/networking
If an IPv6 network address is not specified, this will enable IPv6 internal routing between virtual machines. By default, IPv4 internal routing is enabled.
DNS Domain Name:
Cancel
Back
Finish

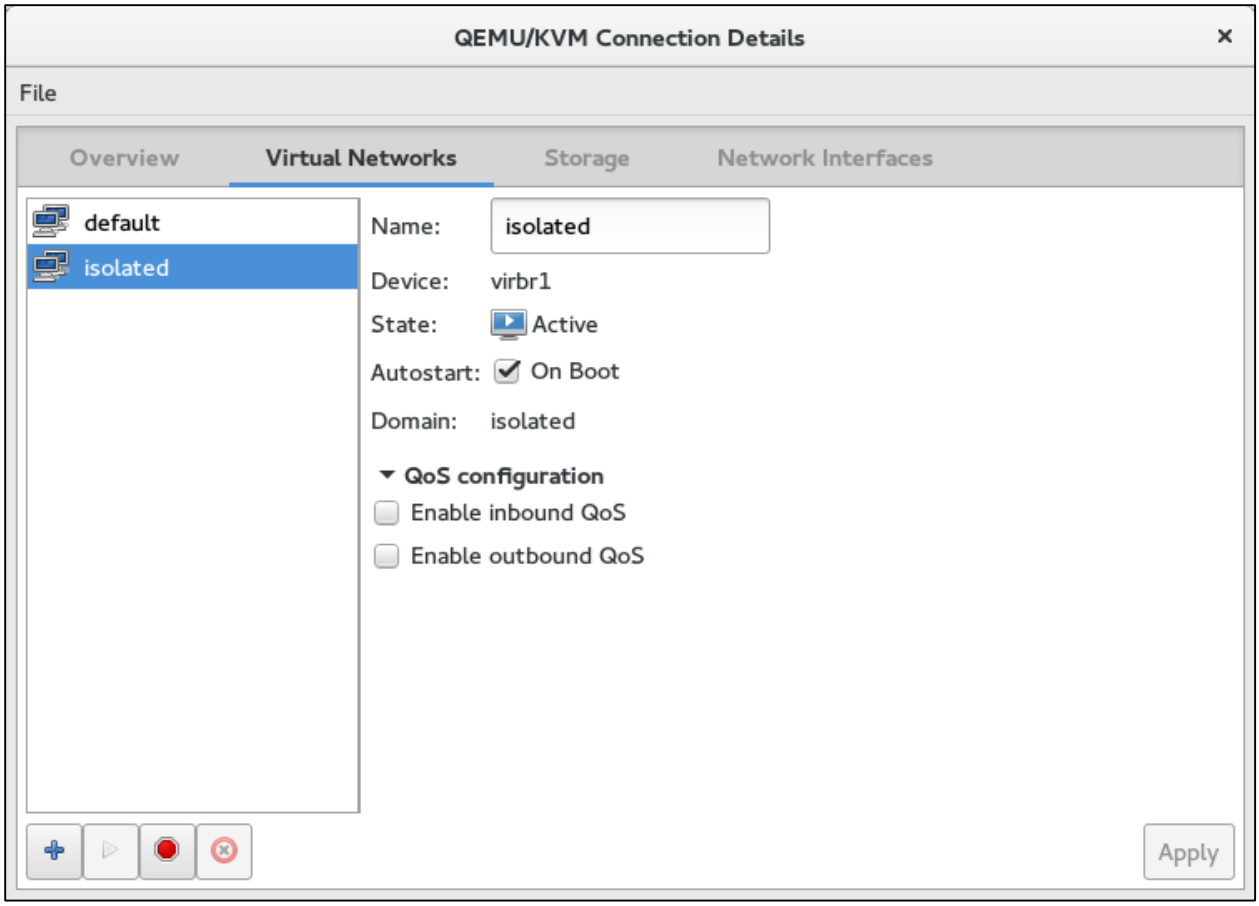
QEMU/KVM Connection Details
File
Overview
Virtual Networks
Storage
Network Interfaces
default
isolated
Name:
isolated
Device:
virbr1
State:
Active
Autostart:
On Boot
Domain:
isolated
QoS configuration
Enable inbound QoS
Enable outbound QoS
Apply

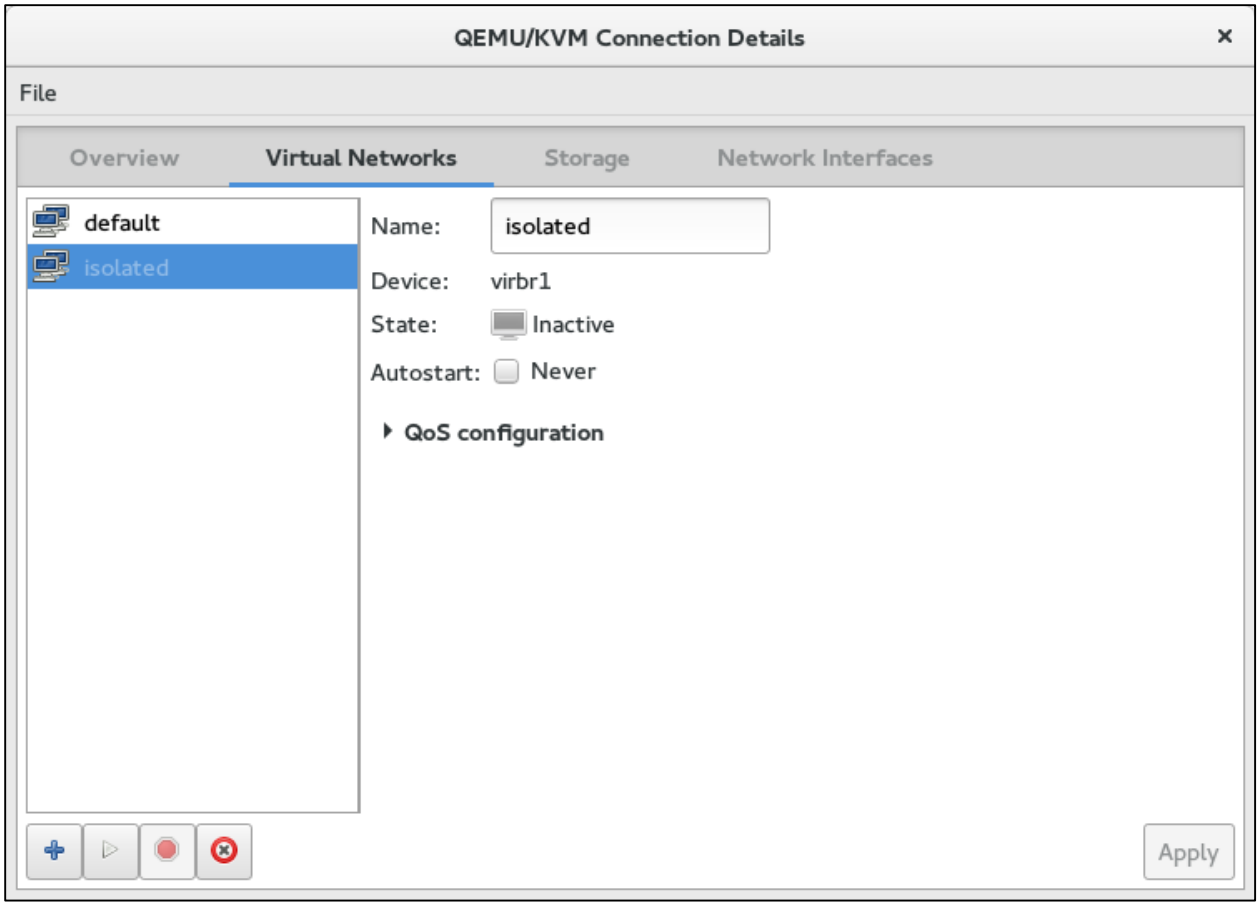
QEMU/KVM Connection Details
File
Overview
Virtual Networks
Storage
Network Interfaces
default
isolated
Name:
isolated
Device:
virbr1
State:
Inactive
Autostart:
Never
QoS configuration
Apply

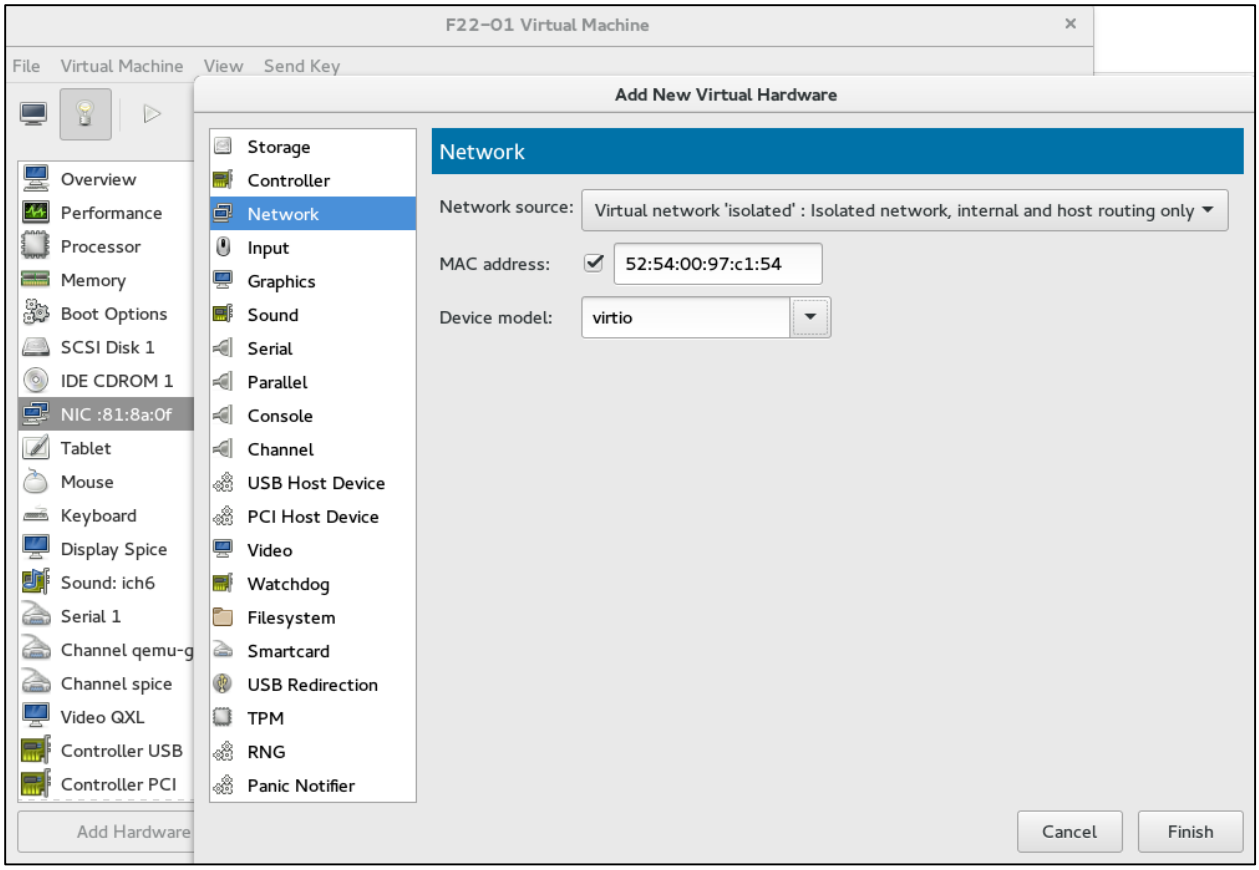
F22-01 Virtual Machine
File
Virtual Machine
View
Send Key
Add New Virtual Hardware
Overview
Performance
Processor
Memory
Boot Options
SCSI Disk 1
IDE CDROM 1
NIC :81:8a:0f
Tablet
Mouse
Keyboard
Display Spice
Sound: ich6
Serial 1
Channel qemu-g
Channel spice
Video QXL
Controller USB
Controller PCI
Add Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Network
Network source:
Virtual network 'isolated' : Isolated network, internal and host routing only
MAC address:
52:54:00:97:c1:54
Device model:
virtio
Cancel
Finish

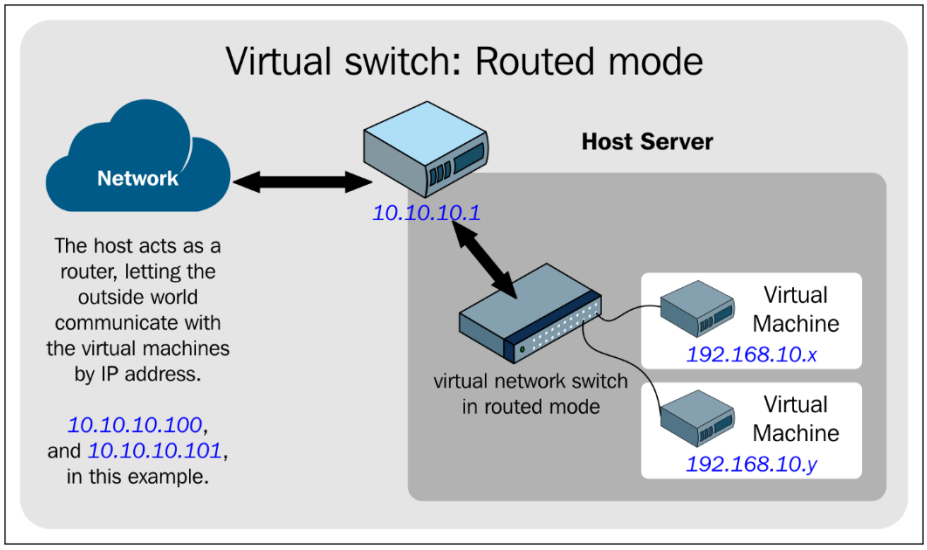
Virtual switch: Routed mode
Network
Host Server
10.10.10.1
The host acts as a router, letting the outside world communicate with the virtual machines by IP address.
10.10.10.100, and 10.10.10.101, in this example.
virtual network switch in routed mode
Virtual Machine
192.168.10.x
Virtual Machine
192.168.10.y

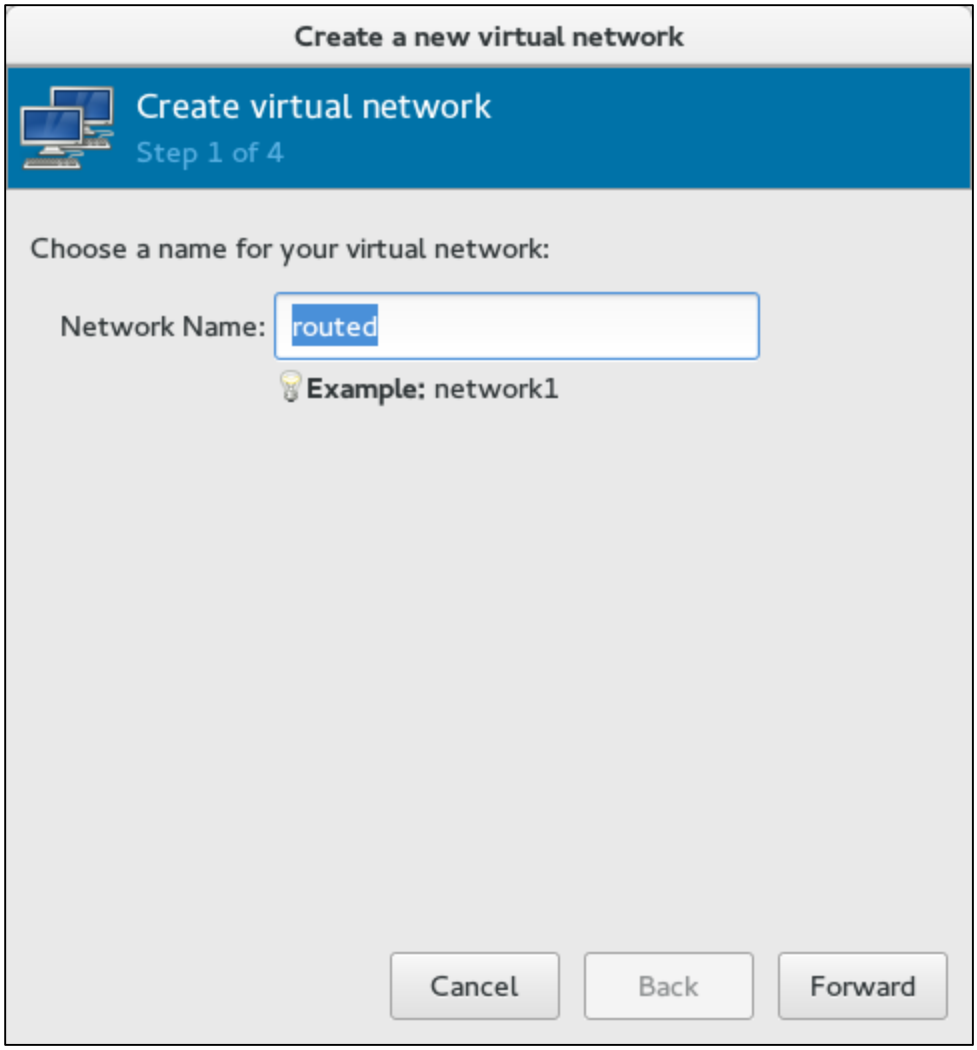
Create a new virtual network
Create virtual network
Step 1 of 4
Choose a name for your virtual network:
Network Name:
routed
Example: network1
Cancel
Back
Forward

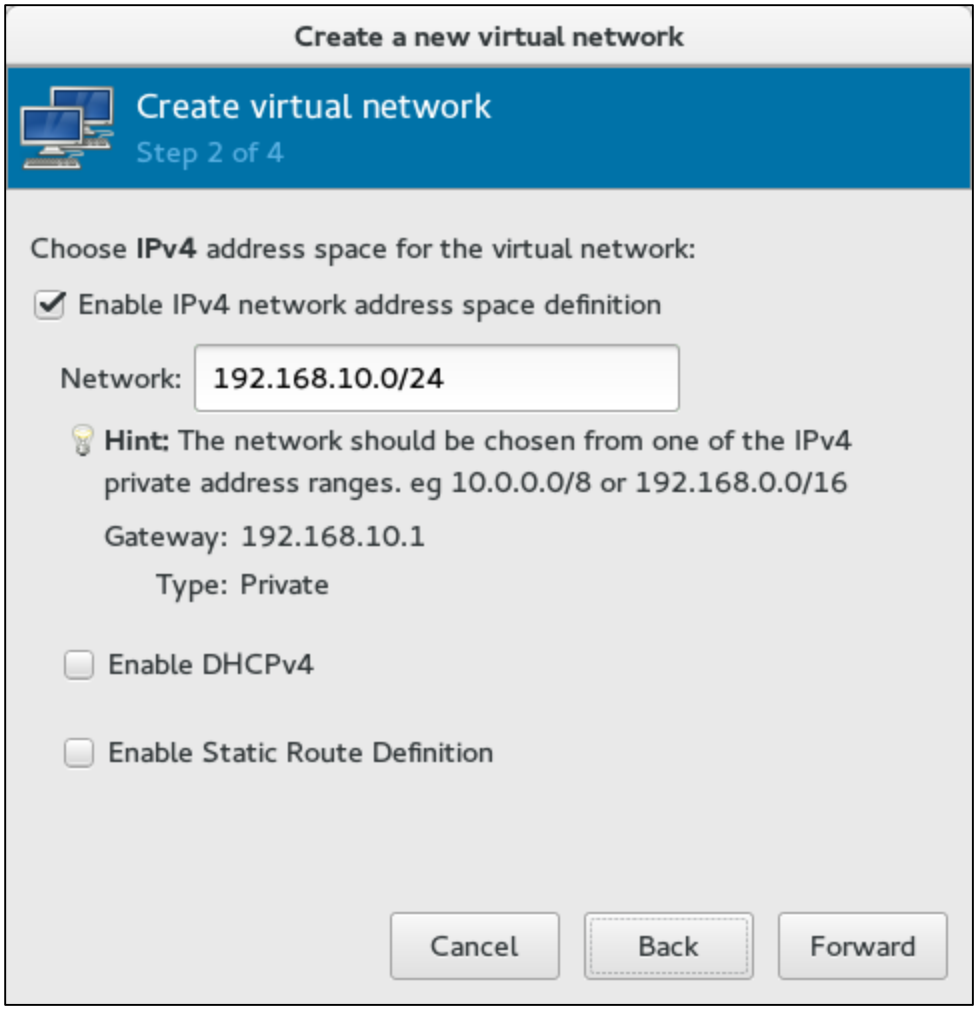
Create a new virtual network
Create virtual network
Step 2 of 4
Choose IPv4 address space for the virtual network:
Enable IPv4 network address space definition
Network:
192.168.10.0/24
Hint: The network should be chosen from one of the IPv4 private address ranges. eg 10.0.0.0/8 or 192.168.0.0/16
Gateway: 192.168.10.1
Type: Private
Enable DHCPv4
Enable Static Route Definition
Cancel
Back
Forward

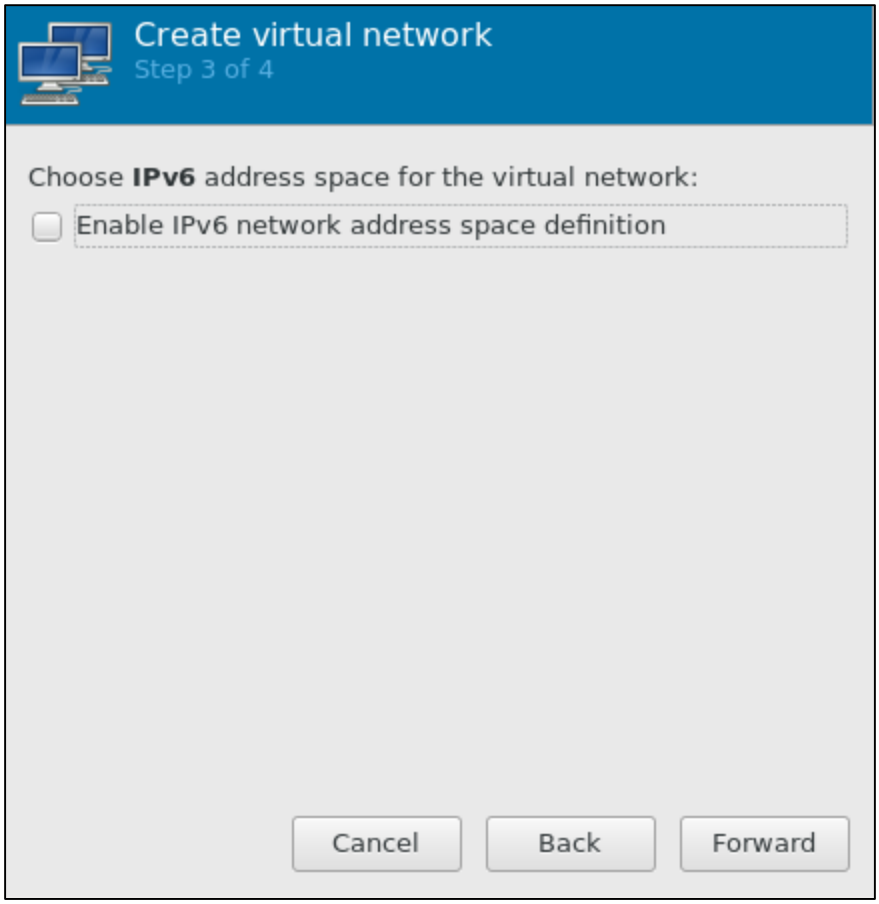
Create virtual network
Step 3 of 4
Choose IPv6 address space for the virtual network:
Enable IPv6 network address space definition
Cancel
Back
Forward

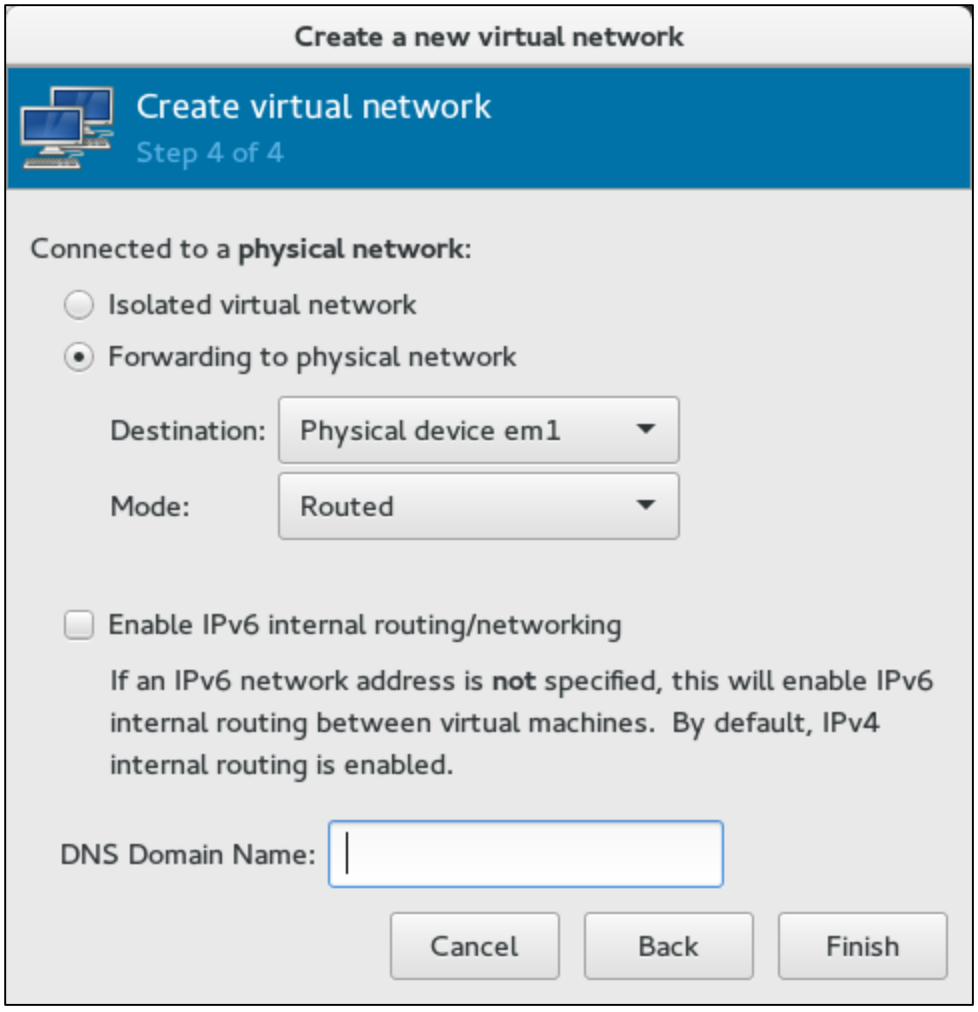


```
<network>
  <name>routed</name>
  <uuid>9a1d8de7-5627-4f08-a3d1-836b7a5fe060</uuid>
  <forward mode='route'/>
  <bridge name='virbr2' stp='on' delay='0'/>
  <mac address='52:54:00:f1:cb:30'/>
  <ip address='192.168.10.1' netmask='255.255.255.0'>
  </ip>
</network>
```

```
error: (network_definition):6: Opening and ending tag mismatch: network line 1 and forward
  </forward>
------------^
Failed. Try again? [y,n,f,?]:
Network routed XML configuration edited.
```

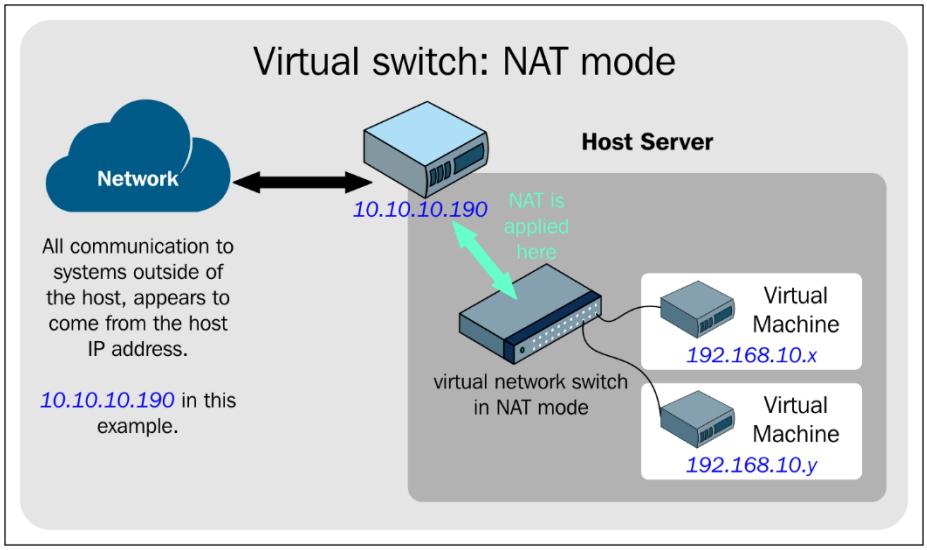
Virtual switch: NAT mode
Network
Host Server
10.10.10.190
NAT is applied here
All communication to systems outside of the host, appears to come from the host IP address.
10.10.10.190 in this example.
virtual network switch in NAT mode
Virtual Machine
192.168.10.x
Virtual Machine
192.168.10.y

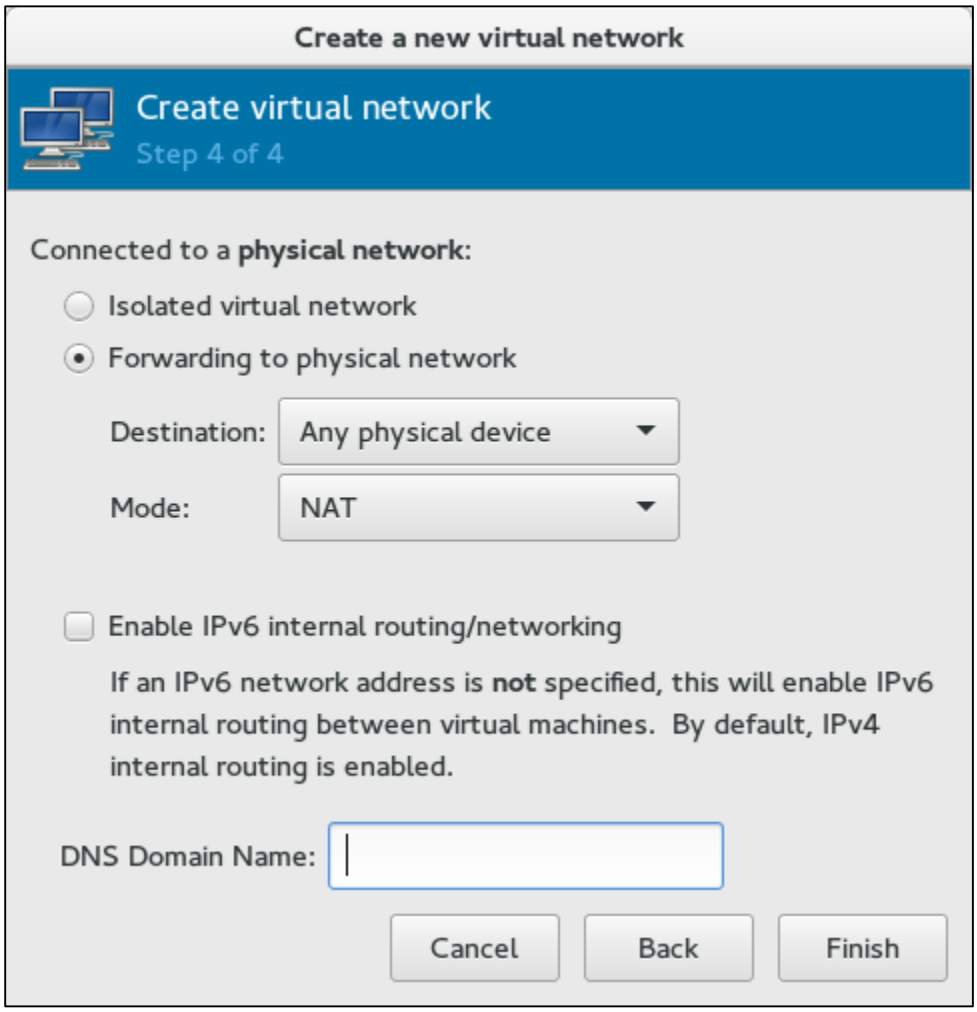
Create a new virtual network
Create virtual network
Step 4 of 4
Connected to a physical network:
Isolated virtual network
Forwarding to physical network
Destination:
Any physical device
Mode:
NAT
Enable IPv6 internal routing/networking
If an IPv6 network address is not specified, this will enable IPv6 internal routing between virtual machines. By default, IPv4 internal routing is enabled.
DNS Domain Name:
Cancel
Back
Finish

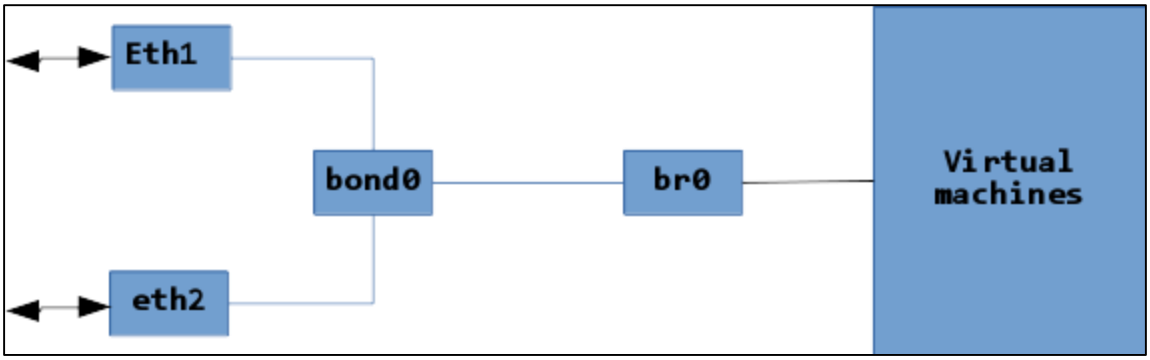
Eth1
bond0
br0
Virtual machines
eth2

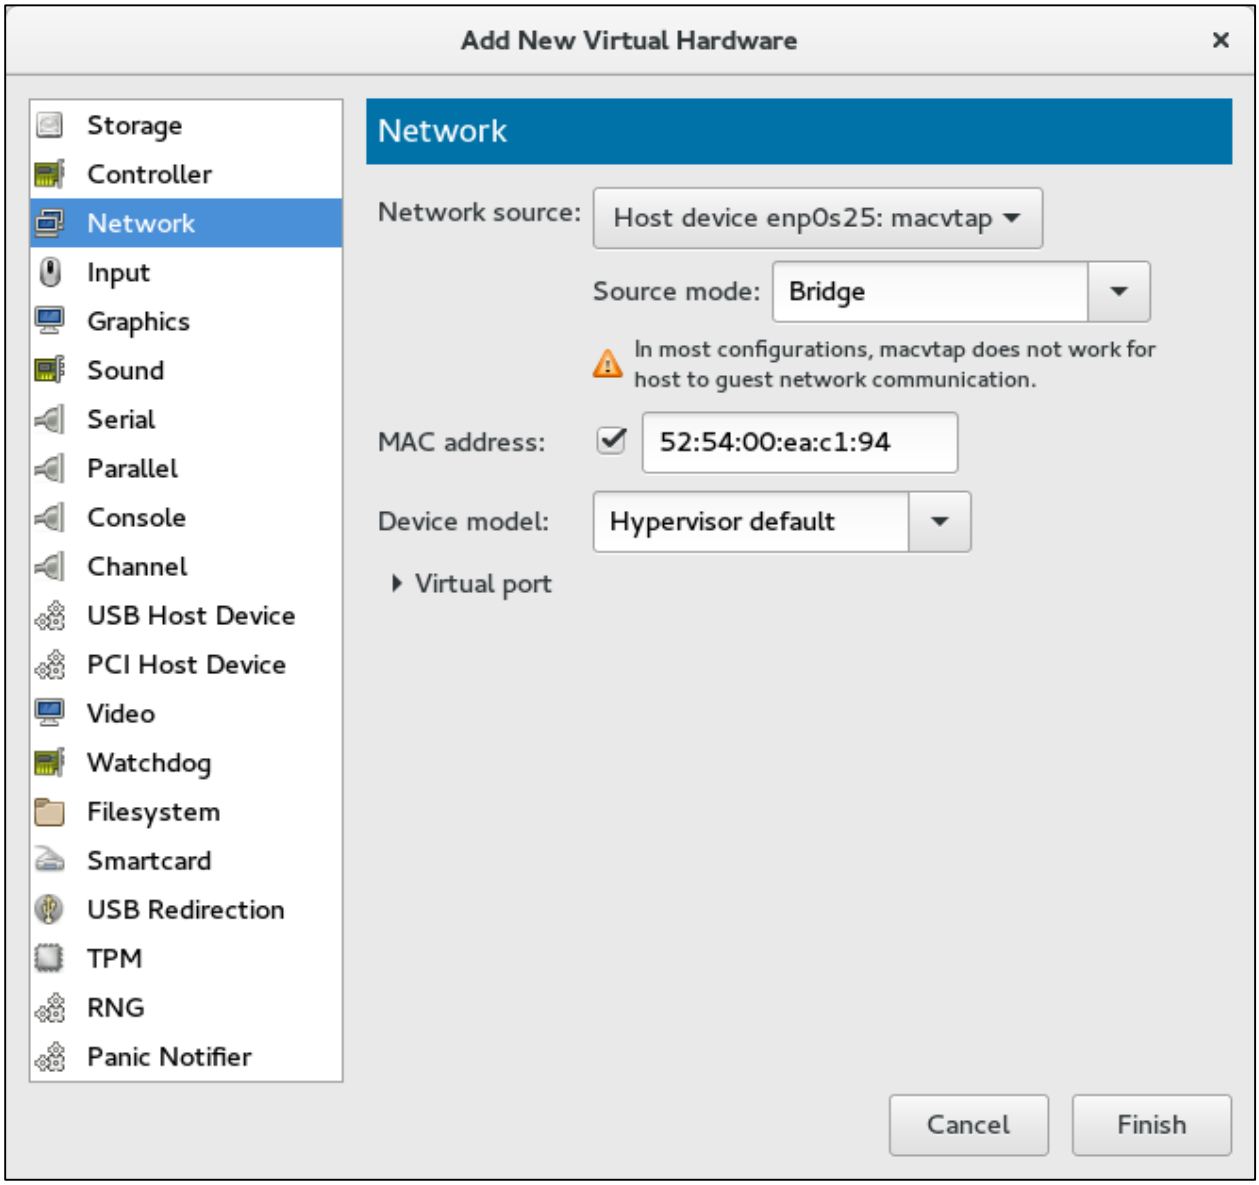
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Network
Network source:
Host device enp0s25: macvtap
Source mode:
Bridge
In most configurations, macvtap does not work for host to guest network communication.
MAC address:
52:54:00:ea:c1:94
Device model:
Hypervisor default
Virtual port
Cancel
Finish


GRUB_CMDLINE_LINUX="
rd.lvm.lv=fedora/swap rhgb quiet intel_iommu=on

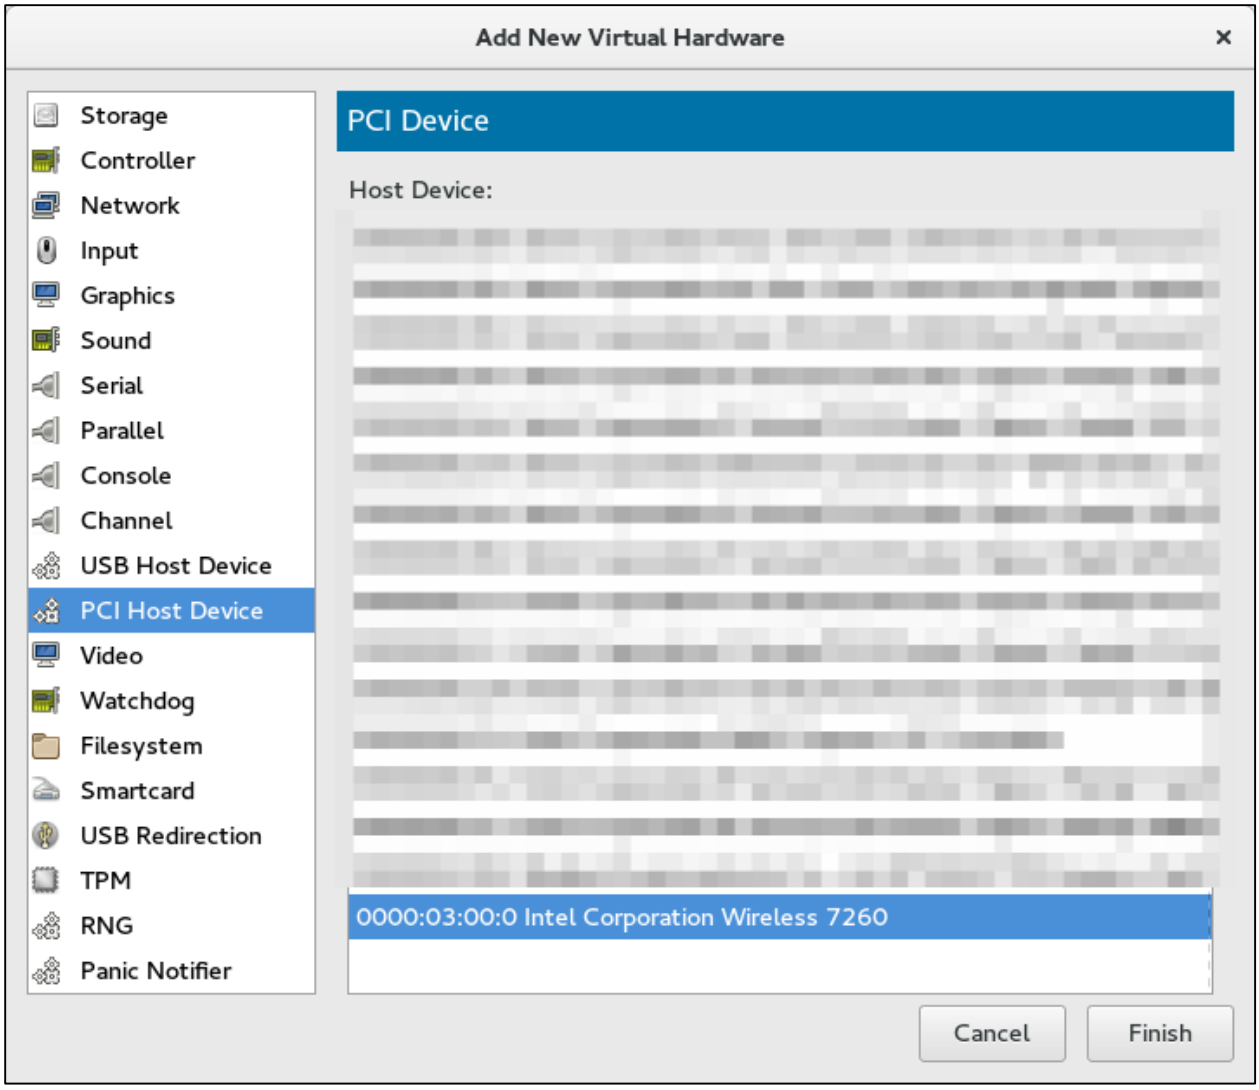
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
PCI Device
Host Device:
0000:03:00:0 Intel Corporation Wireless 7260
Cancel
Finish

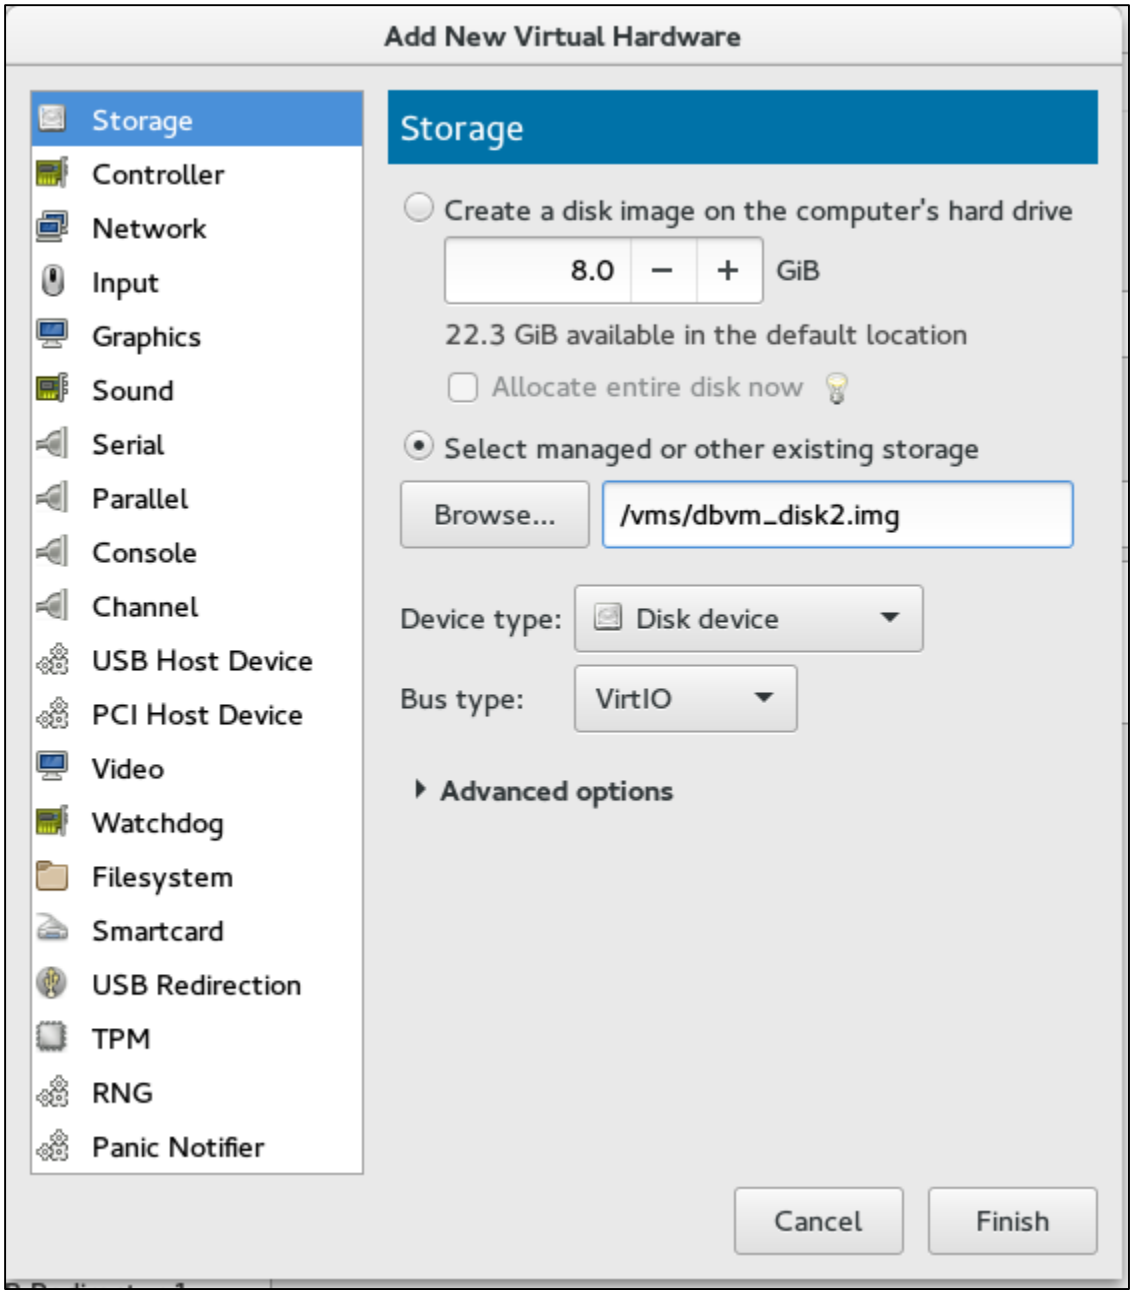
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Storage
Create a disk image on the computer's hard drive
8.0
GiB
22.3 GiB available in the default location
Allocate entire disk now
Select managed or other existing storage
Browse...
/vms/dbvm_disk2.img
Device type:
Disk device
Bus type:
VirtIO
Advanced options
Cancel
Finish

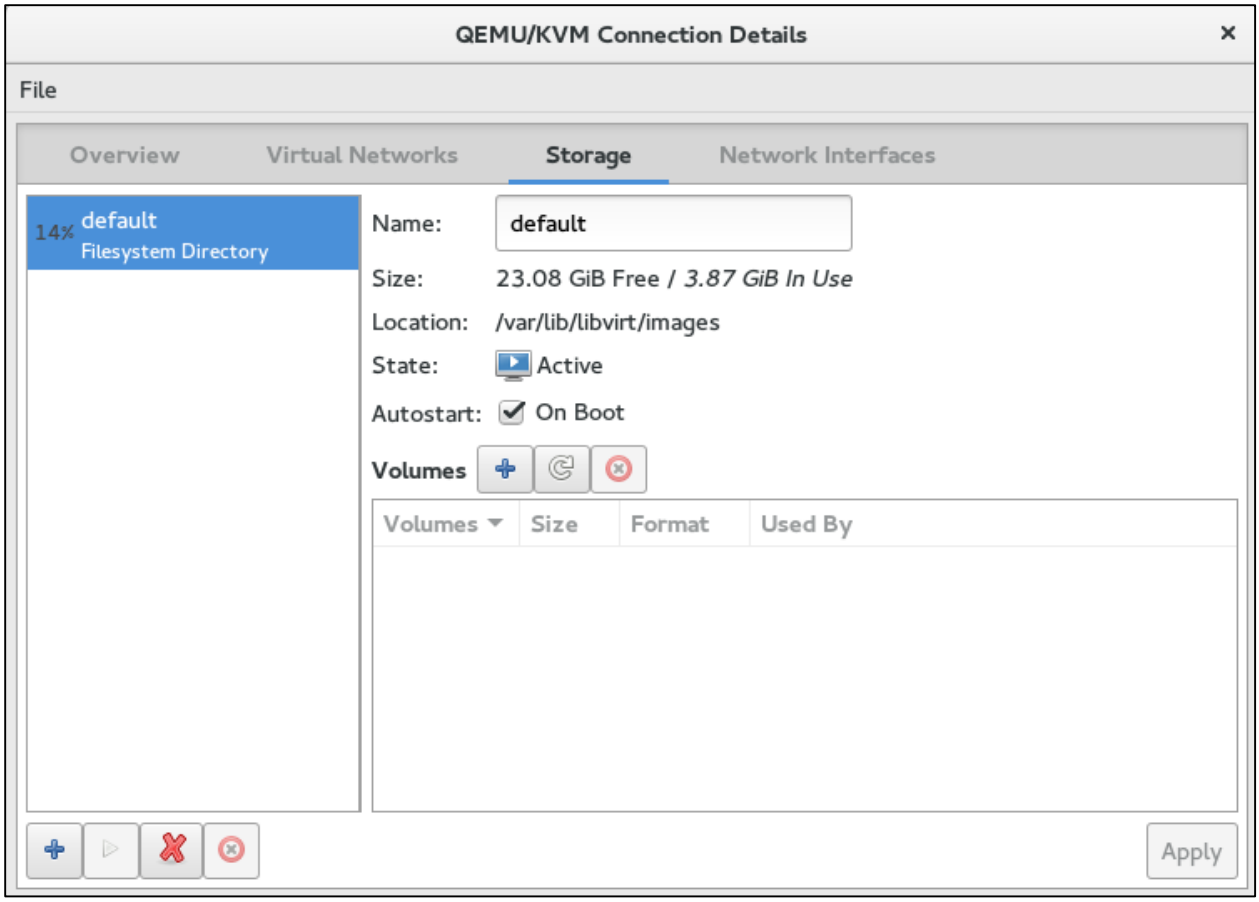
QEMU/KVM Connection Details
File
Overview
Virtual Networks
Storage
Network Interfaces
14% default
Filesystem Directory
Name:
default
Size:
23.08 GiB Free / 3.87 GiB In Use
Location:
/var/lib/libvirt/images
State:
Active
Autostart:
On Boot
Volumes
Volumes
Size
Format
Used By
Apply

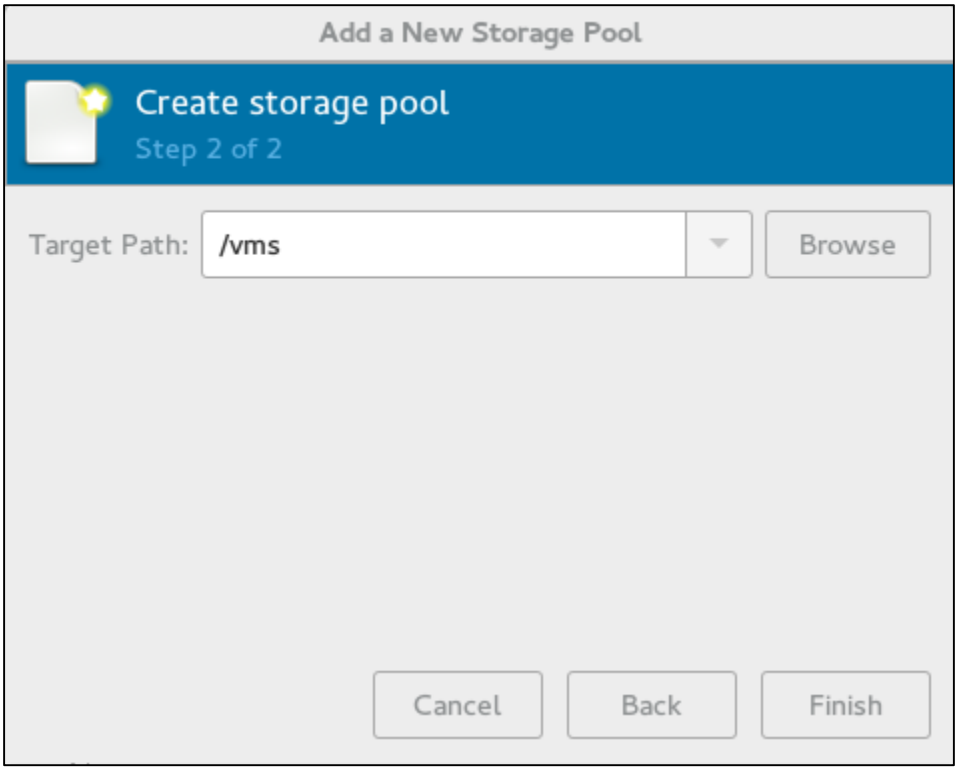
Add a New Storage Pool
Create storage pool
Step 2 of 2
Target Path:
/vms
Browse
Cancel
Back
Finish

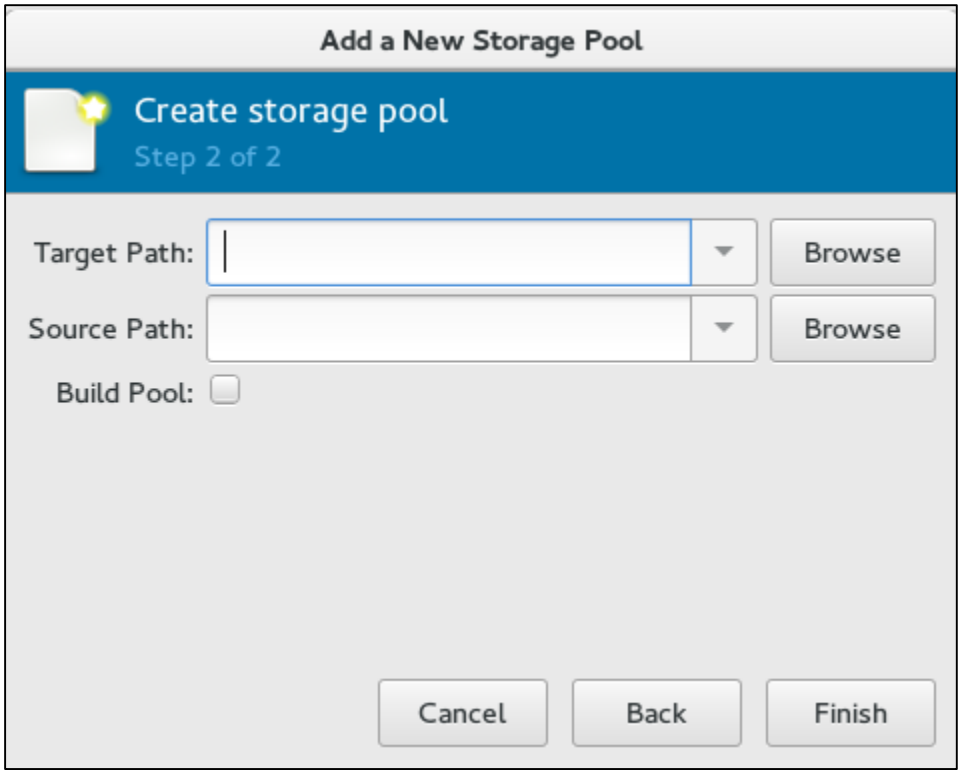
Add a New Storage Pool
Create storage pool
Step 2 of 2
Target Path:
Browse
Source Path:
Browse
Build Pool:
Cancel
Back
Finish

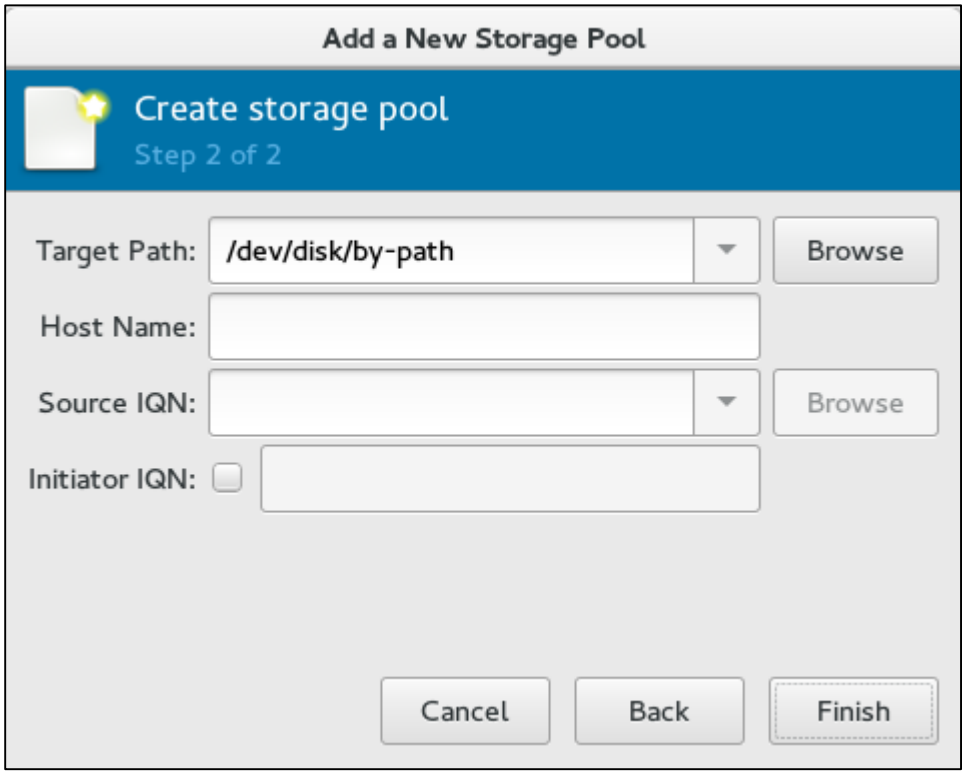
Add a New Storage Pool
Create storage pool
Step 2 of 2
Target Path:
/dev/disk/by-path
Browse
Host Name:
Source IQN:
Browse
Initiator IQN:
Cancel
Back
Finish

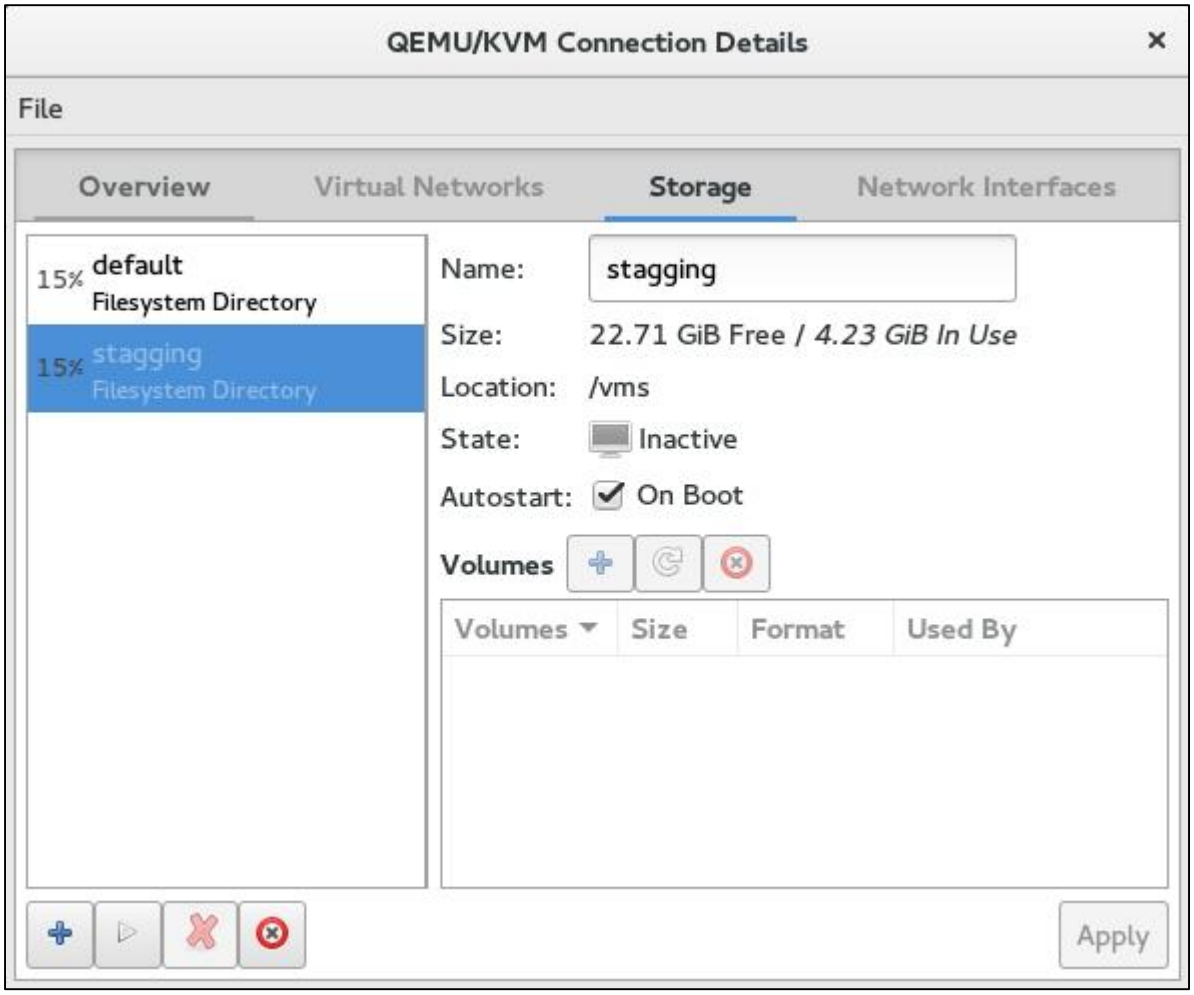
QEMU/KVM Connection Details
File
Overview
Virtual Networks
Storage
Network Interfaces
15% default
Filesystem Directory
15% stagging
Filesystem Directory
Name:
stagging
Size:
22.71 GiB Free / 4.23 GiB In Use
Location:
/vms
State:
Inactive
Autostart:
On Boot
Volumes
Volumes
Size
Format
Used By
Apply

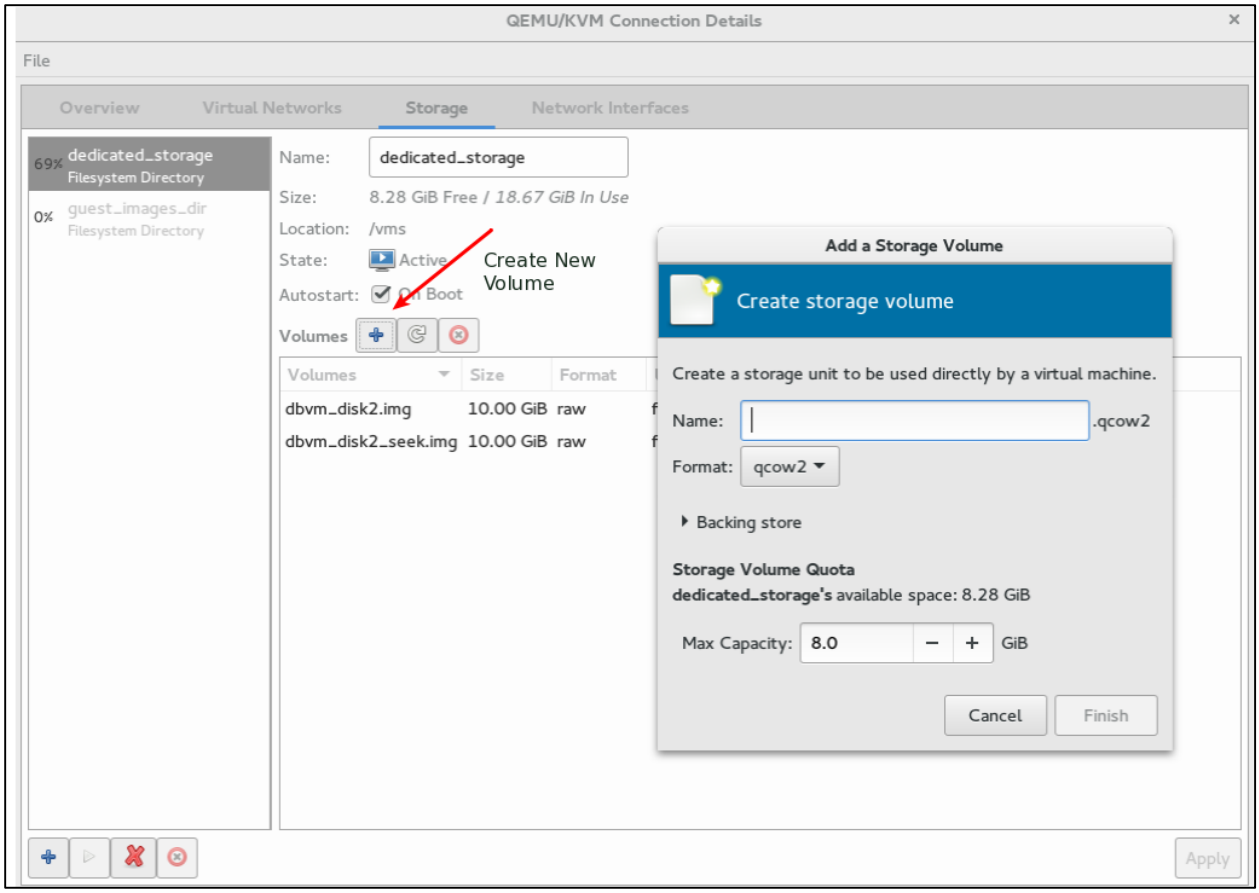
QEMU/KVM Connection Details
File
Overview
Virtual Networks
Storage
Network Interfaces
69% dedicated_storage
Filesystem Directory
0% guest_images_dir
Filesystem Directory
Name: dedicated_storage
Size: 8.28 GiB Free / 18.67 GiB In Use
Location: /vms
State: Active
Autostart: On Boot
Volumes
Create New Volume
Volumes
Size
Format
dbvm_disk2.img 10.00 GiB raw
dbvm_disk2_seek.img 10.00 GiB raw
Add a Storage Volume
Create storage volume
Create a storage unit to be used directly by a virtual machine.
Name: .qcow2
Format: qcow2
Backing store
Storage Volume Quota
dedicated_storage's available space: 8.28 GiB
Max Capacity: 8.0 GiB
Cancel
Finish
Apply

Chapter 6: Virtual Machine Lifecycle Management

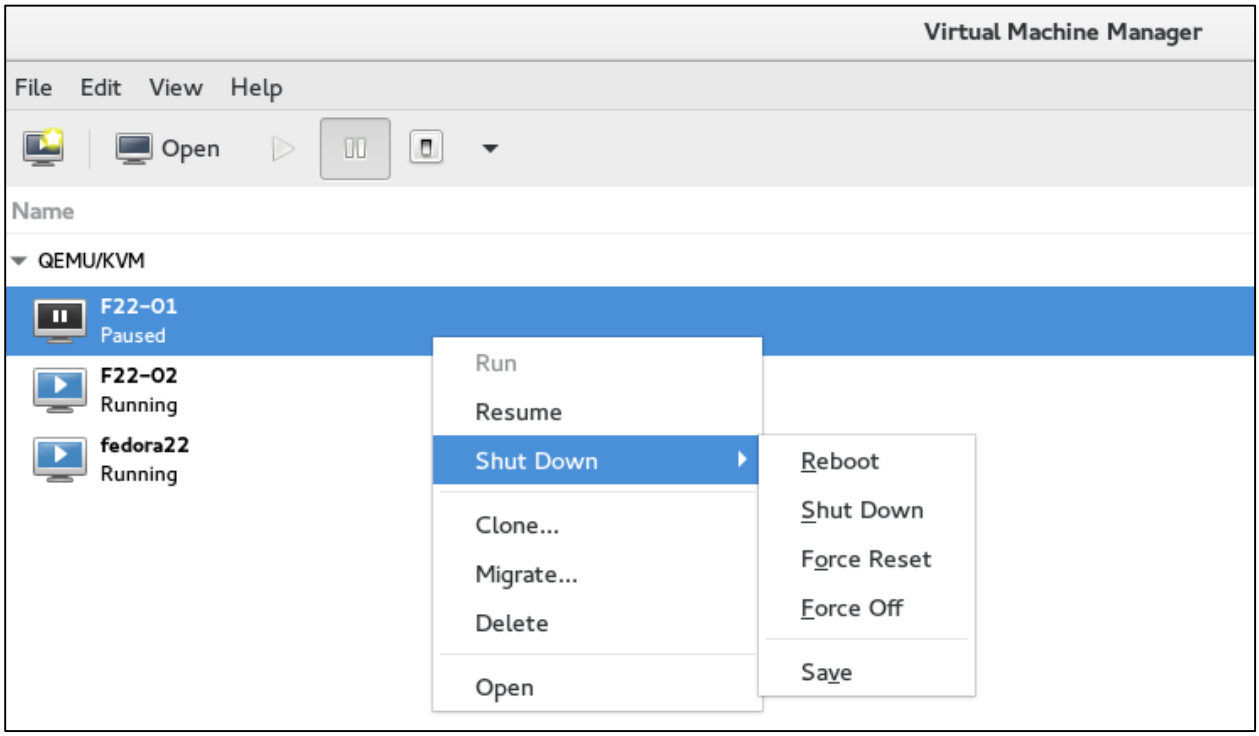

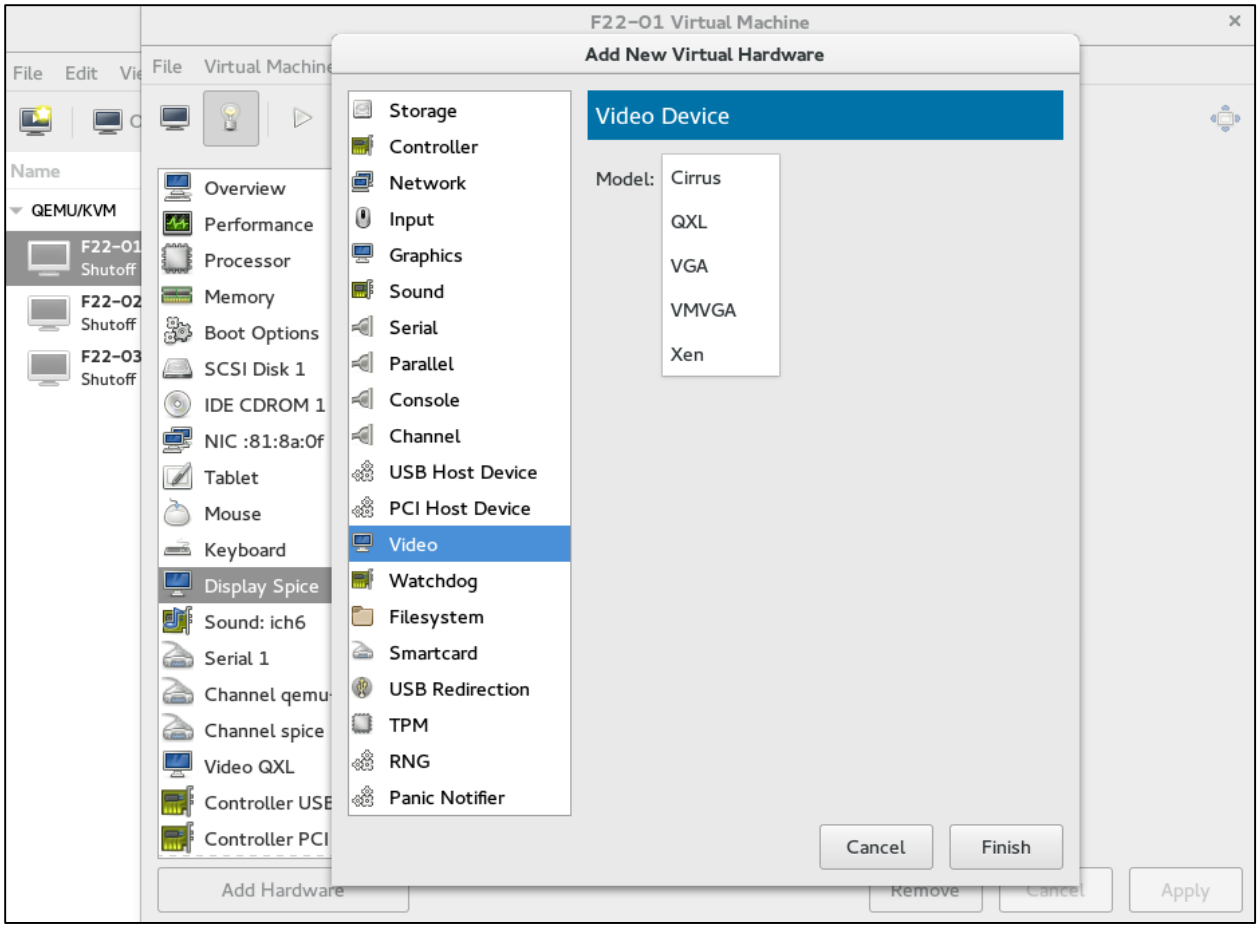
F22-01 Virtual Machine
File Edit Vie
File Virtual Machine
Name
QEMU/KVM
F22-01
Shutoff
F22-02
Shutoff
F22-03
Shutoff
Overview
Performance
Processor
Memory
Boot Options
SCSI Disk 1
IDE CDROM 1
NIC :81:8a:0f
Tablet
Mouse
Keyboard
Display Spice
Sound: ich6
Serial 1
Channel qemu
Channel spice
Video QXL
Controller USB
Controller PCI
Add Hardware
Remove
Cancel
Apply
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Video Device
Model:
Cirrus
QXL
VGA
VMVGA
Xen
Cancel
Finish

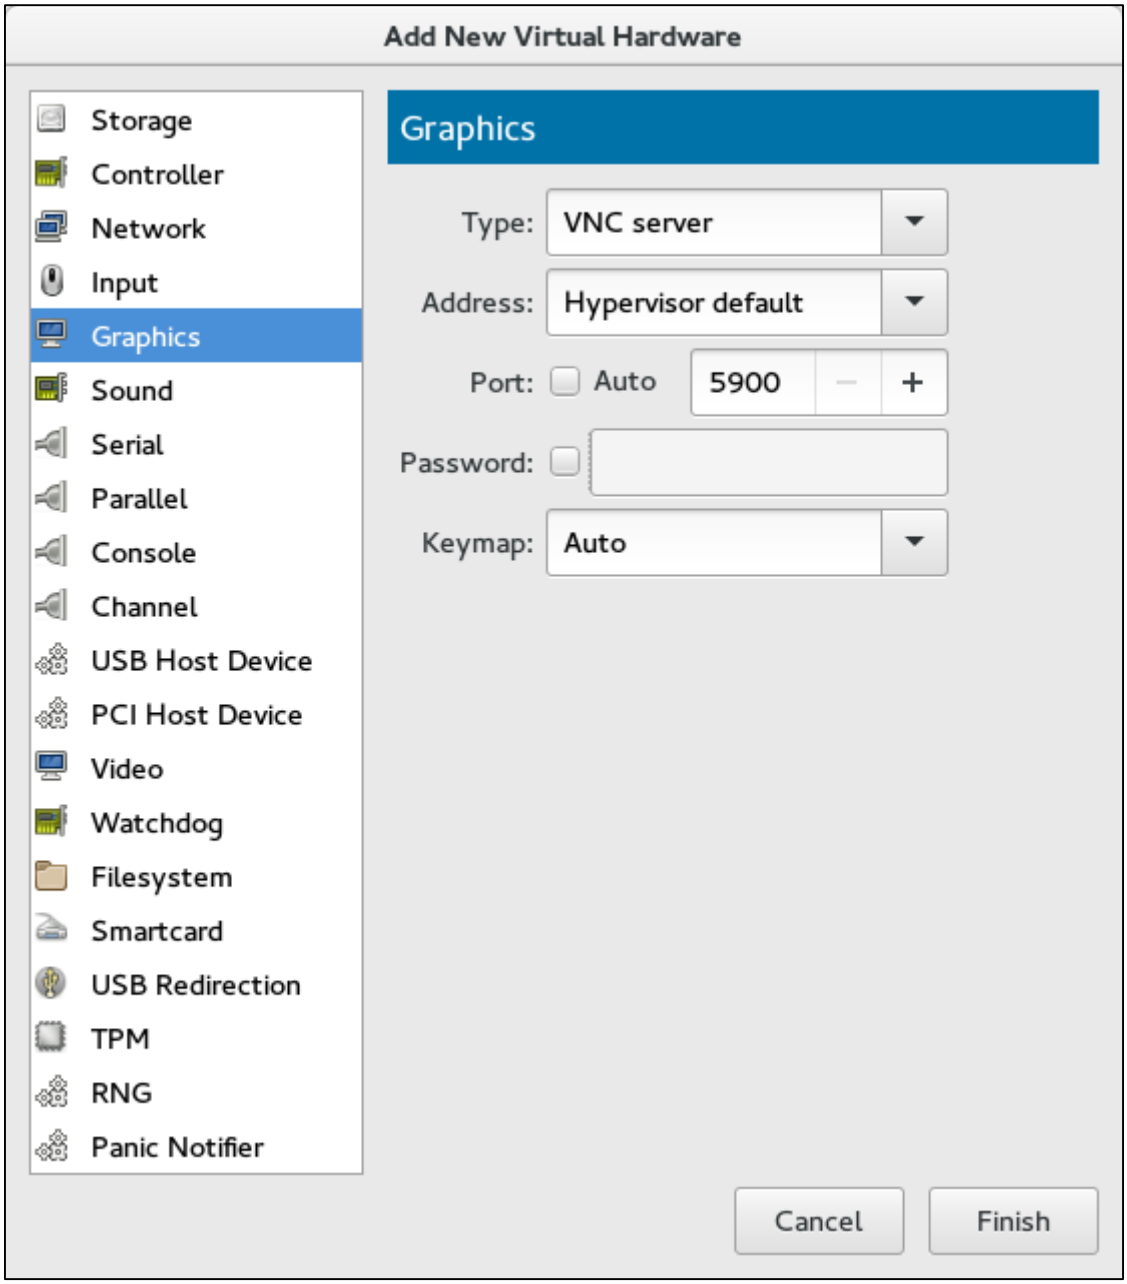
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Graphics
Type:
VNC server
Address:
Hypervisor default
Port:
Auto
5900
Password:
Keymap:
Auto
Cancel
Finish

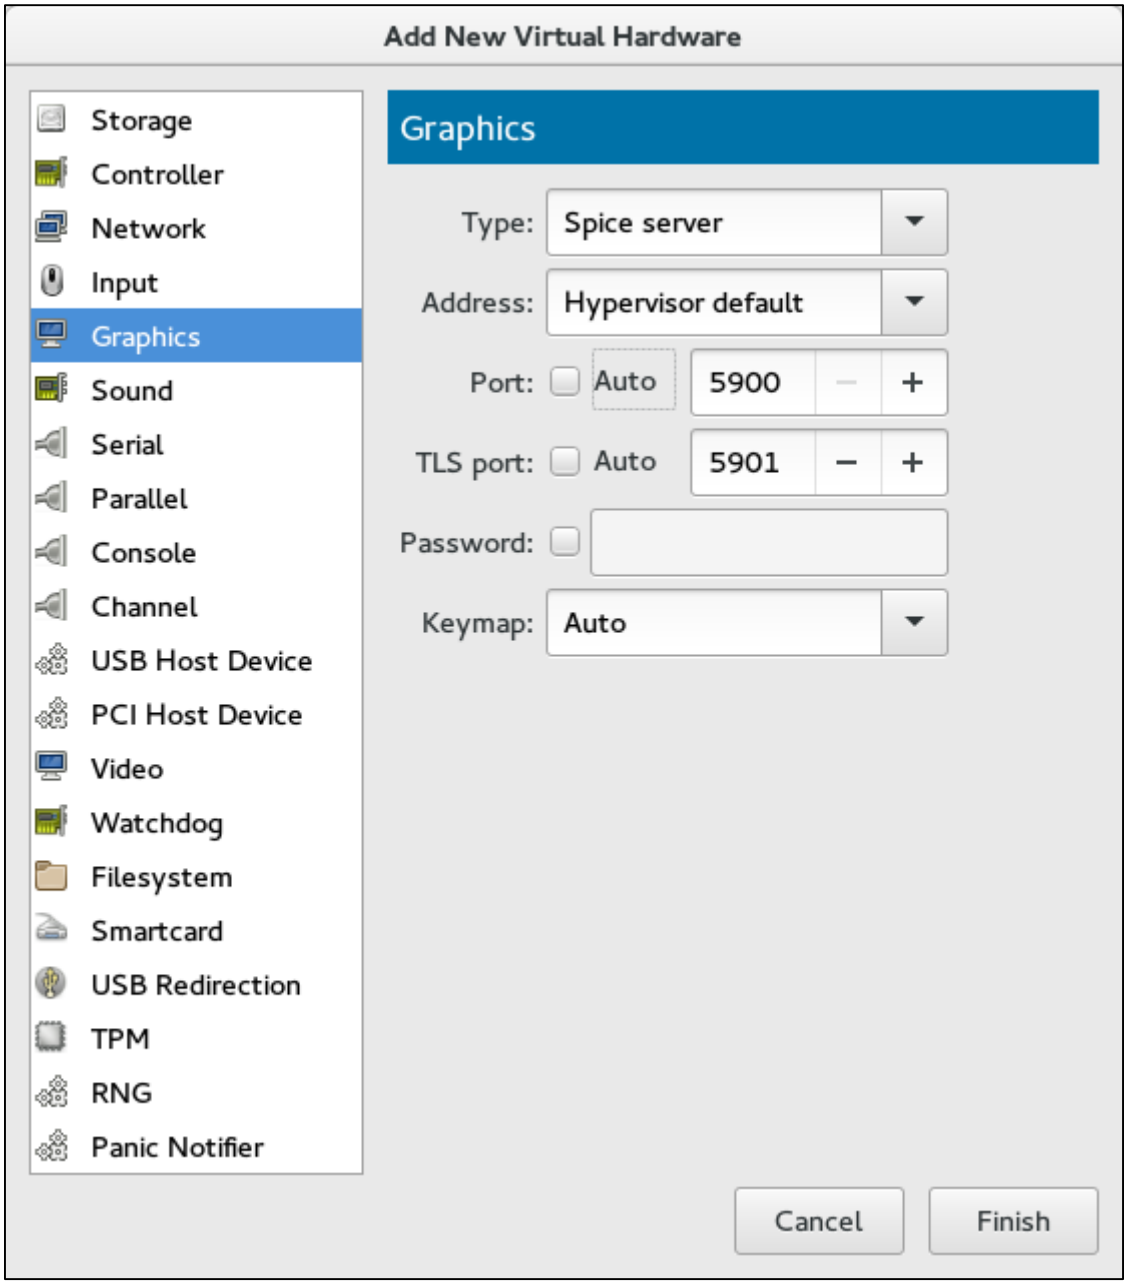
Add New Virtual Hardware
Storage
Controller
Network
Input
Graphics
Sound
Serial
Parallel
Console
Channel
USB Host Device
PCI Host Device
Video
Watchdog
Filesystem
Smartcard
USB Redirection
TPM
RNG
Panic Notifier
Graphics
Type:
Spice server
Address:
Hypervisor default
Port:
Auto
5900
TLS port:
Auto
5901
Password:
Keymap:
Auto
Cancel
Finish

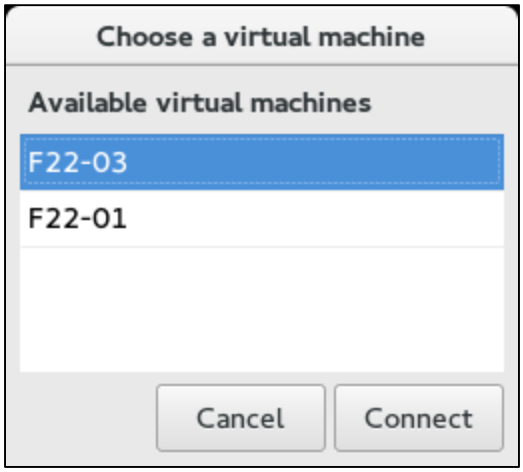
Choose a virtual machine
Available virtual machines
F22-03
F22-01
Cancel
Connect

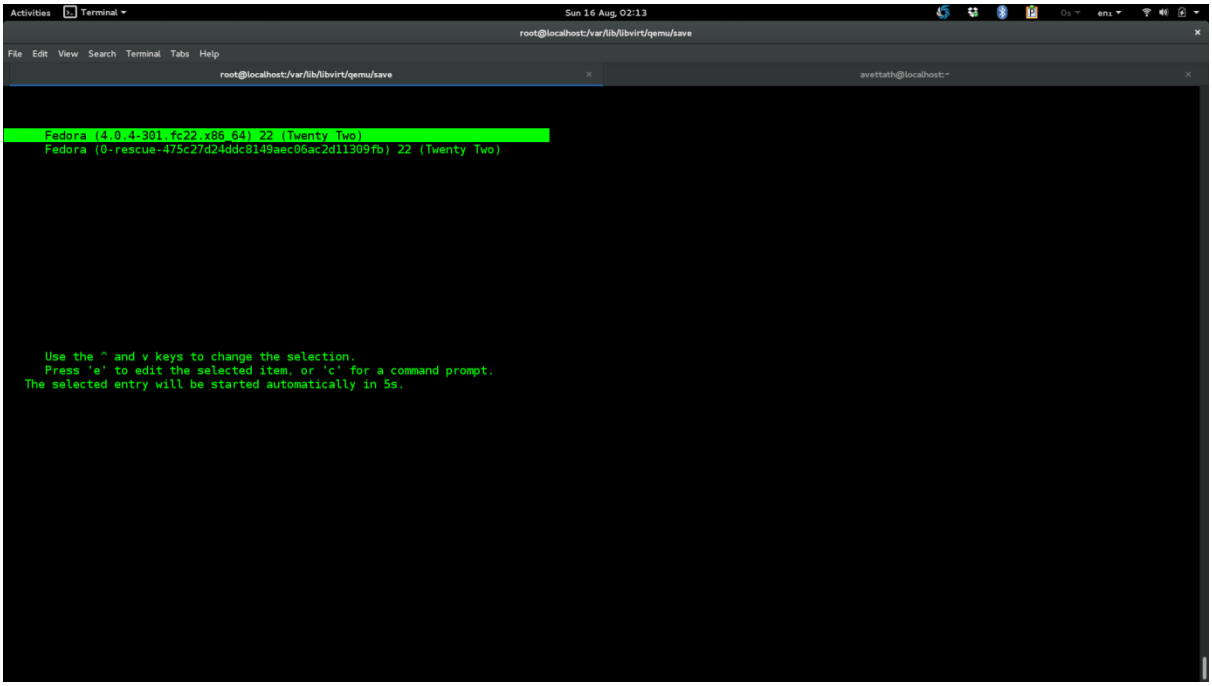
Activities
Terminal
Sun 16 Aug, 02:13
root@localhost:/var/lib/libvirt/qemu/save
File Edit View Search Terminal Tabs Help
root@localhost:/var/lib/libvirt/qemu/save
avettath@localhost:~
Fedora (4.0.4-301.fc22.x86_64) 22 (Twenty Two)
Fedora (0-rescue-475c27d24ddc8149aec06ac2d11309fb) 22 (Twenty Two)
Use the ^ and v keys to change the selection.
Press 'e' to edit the selected item, or 'c' for a command prompt.
The selected entry will be started automatically in 5s.

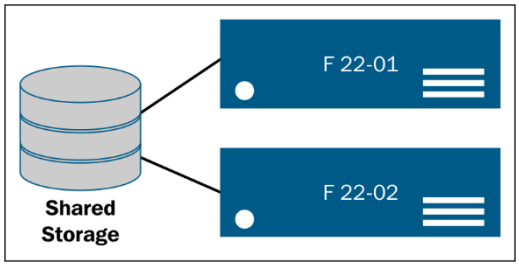
F 22-01
F 22-02
Shared
Storage

Chapter 7: Templates and Snapshots

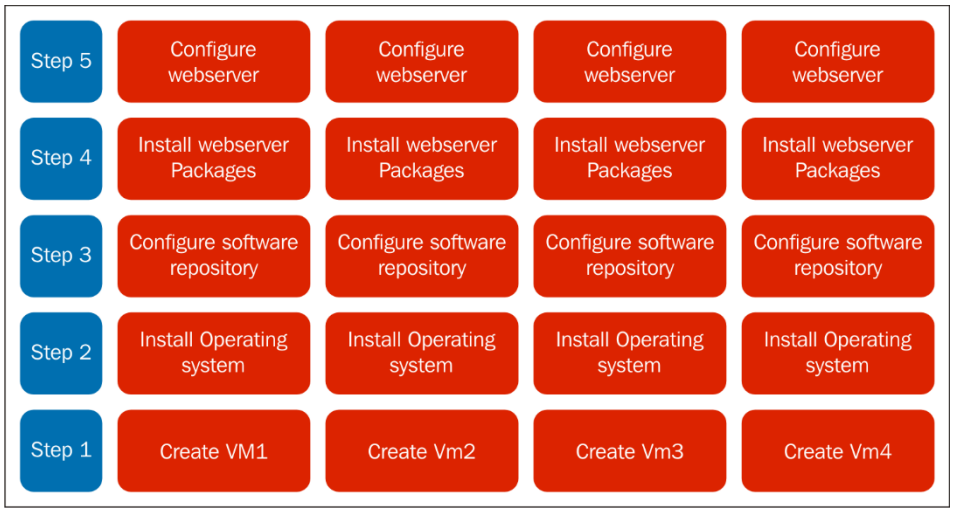

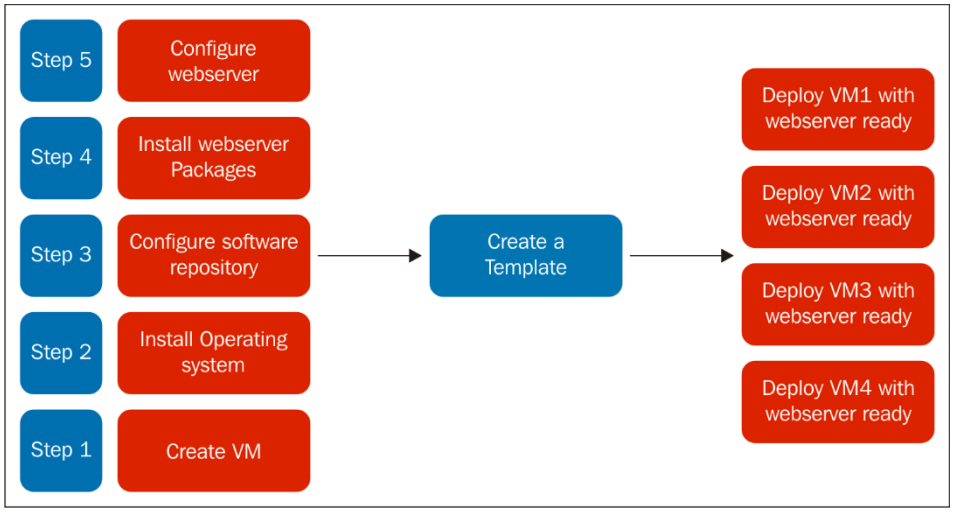

Clone Virtual Machine

Clone virtual machine

Create clone based on: Template_CentOS

Name: Template_CentOS-clone

Networking: NAT (52:54:00:fe:4d:ff) Details...

Storage: CentOS7

Clone this disk (8.0 GiB)

CentOS-7-x86_64-Minimal-1503-01.iso (Removable, Read Only)

Share disk with Template_CentOS

Cloning creates a new, independent copy of the original disk. Sharing uses the existing disk image for both the original and the new machine.

Cloning does not alter the guest OS contents. If you need to do things like change passwords or static IPs, please see the virt-sysprep(1) tool.

Cancel Clone

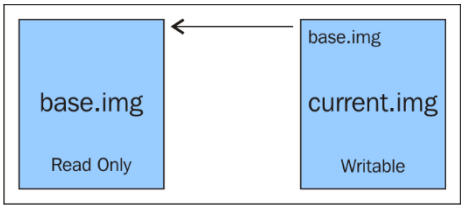

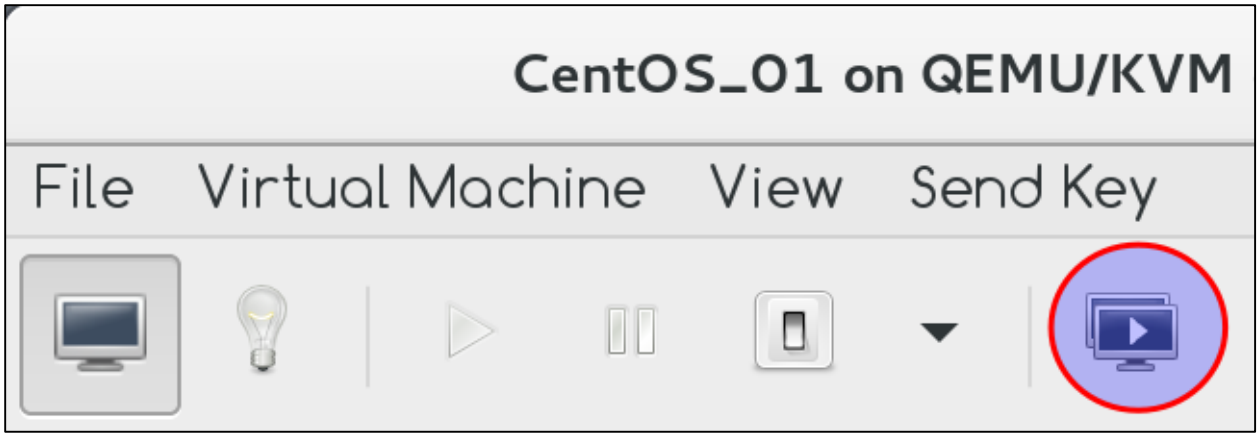
CentOS_01 on QEMU/KVM
File
Virtual Machine
View
Send Key

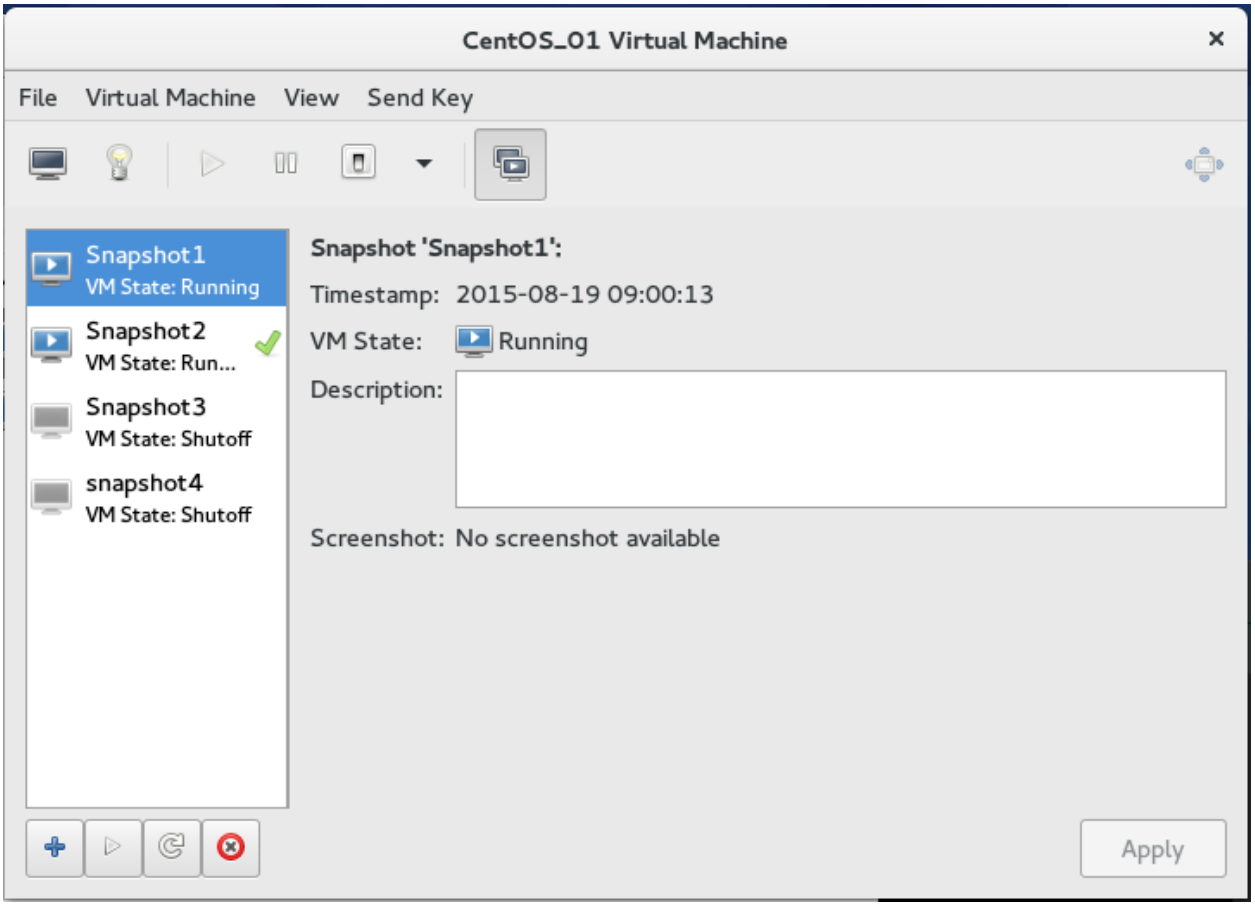
CentOS_01 Virtual Machine
File
Virtual Machine
View
Send Key
Snapshot1
VM State: Running
Snapshot2
VM State: Run...
Snapshot3
VM State: Shutoff
snapshot4
VM State: Shutoff
Snapshot 'Snapshot1':
Timestamp: 2015-08-19 09:00:13
VM State: Running
Description:
Screenshot: No screenshot available
Apply

Create snapshot

Create snapshot

Name: snapshot2

Status: Running

Description: Snapshot prior applying patch set #104

Screenshot:

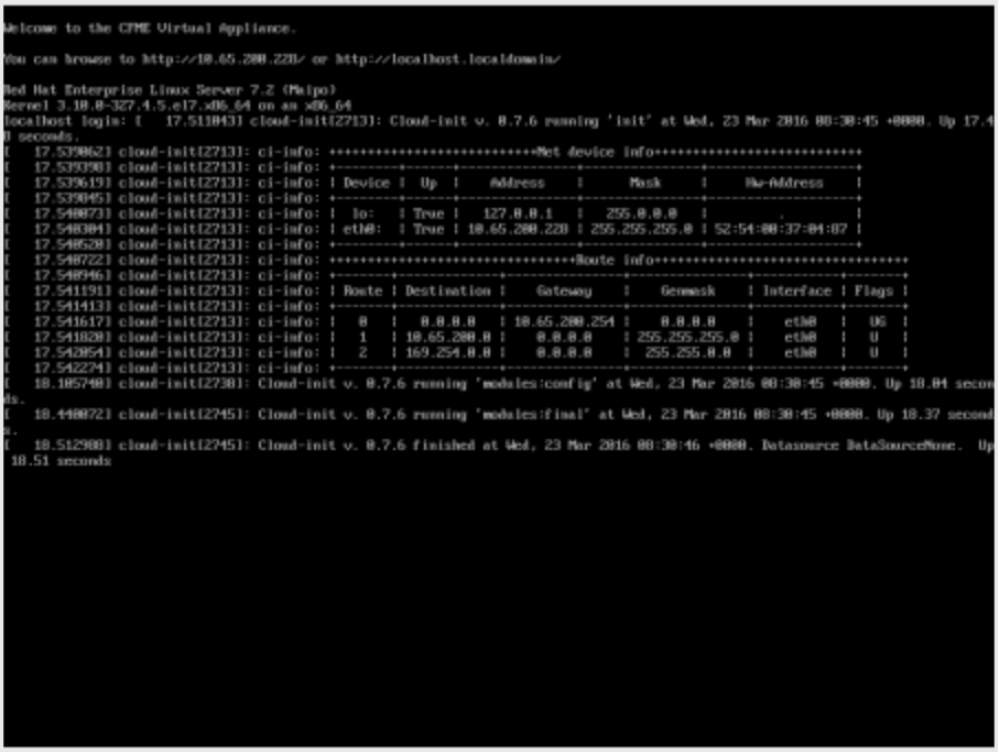

Cancel Finish

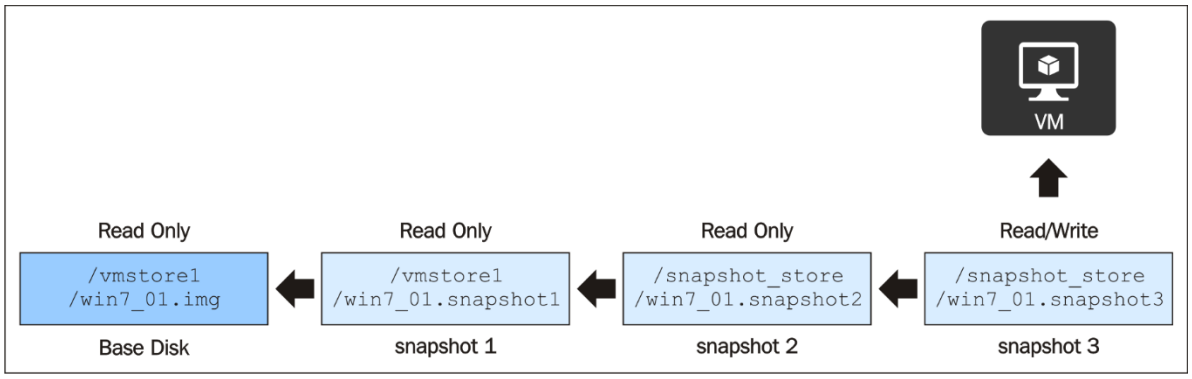
VM
Read Only
/vmstore1
/win7_01.img
Base Disk
Read Only
/vmstore1
/win7_01.snapshot1
snapshot 1
Read Only
/snapshot_store
/win7_01.snapshot2
snapshot 2
Read/Write
/snapshot_store
/win7_01.snapshot3
snapshot 3

Chapter 8: Kimchi – An HTML5-Based Management Tool for KVM/libvirt

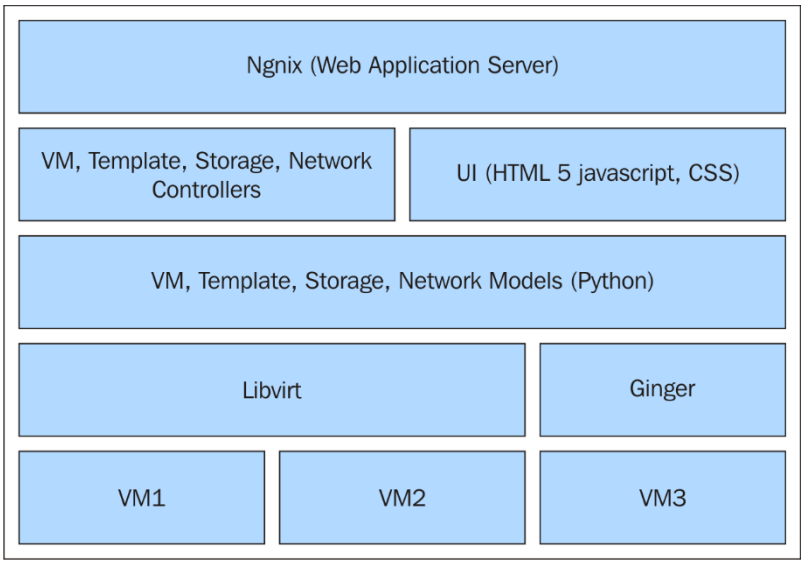

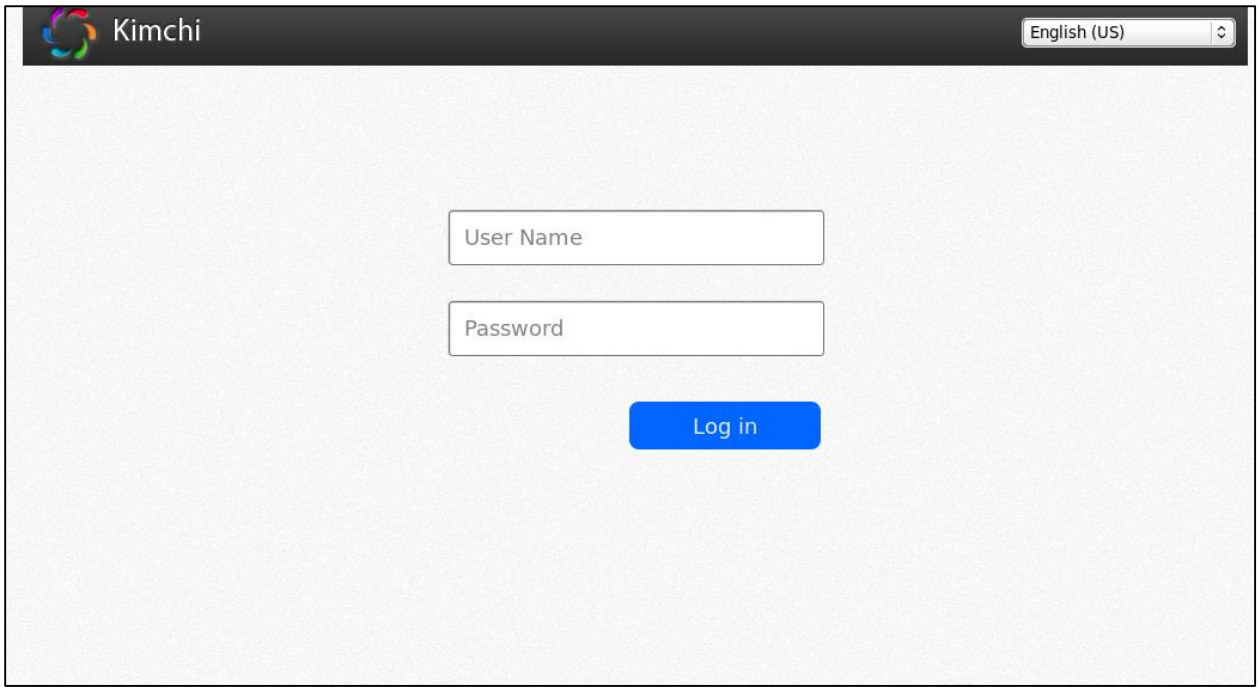

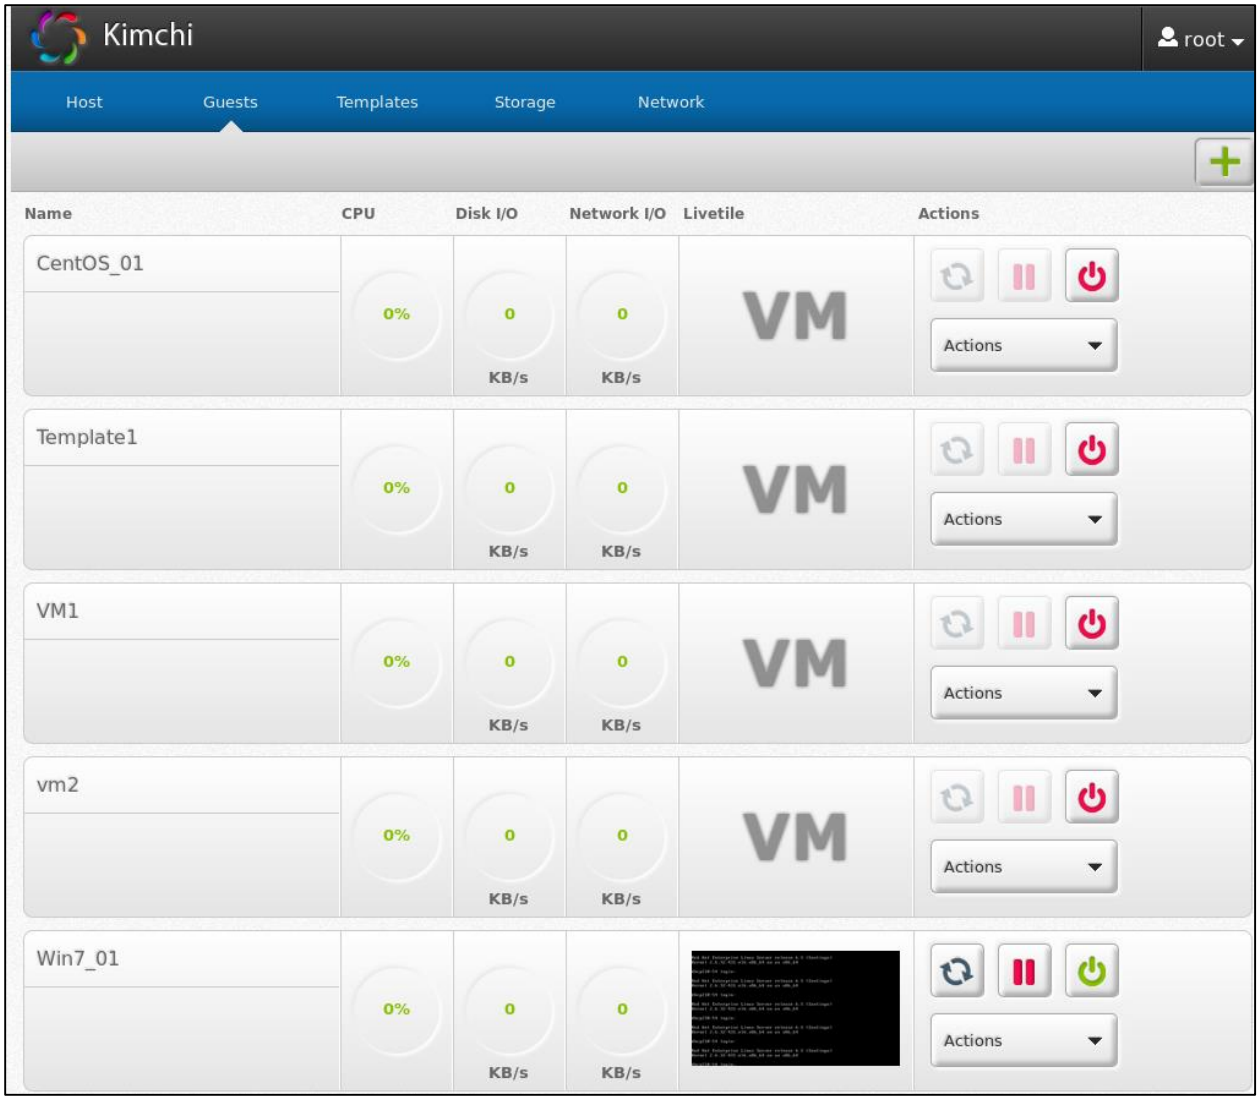
Kimchi
root
Host
Guests
Templates
Storage
Network
Name
CPU
Disk I/O
Network I/O
Livetile
Actions
CentOS_01
0%
0
KB/s
0
KB/s
VM
Actions
Template1
0%
0
KB/s
0
KB/s
VM
Actions
VM1
0%
0
KB/s
0
KB/s
VM
Actions
vm2
0%
0
KB/s
0
KB/s
VM
Actions
Win7_01
0%
0
KB/s
0
KB/s
Actions

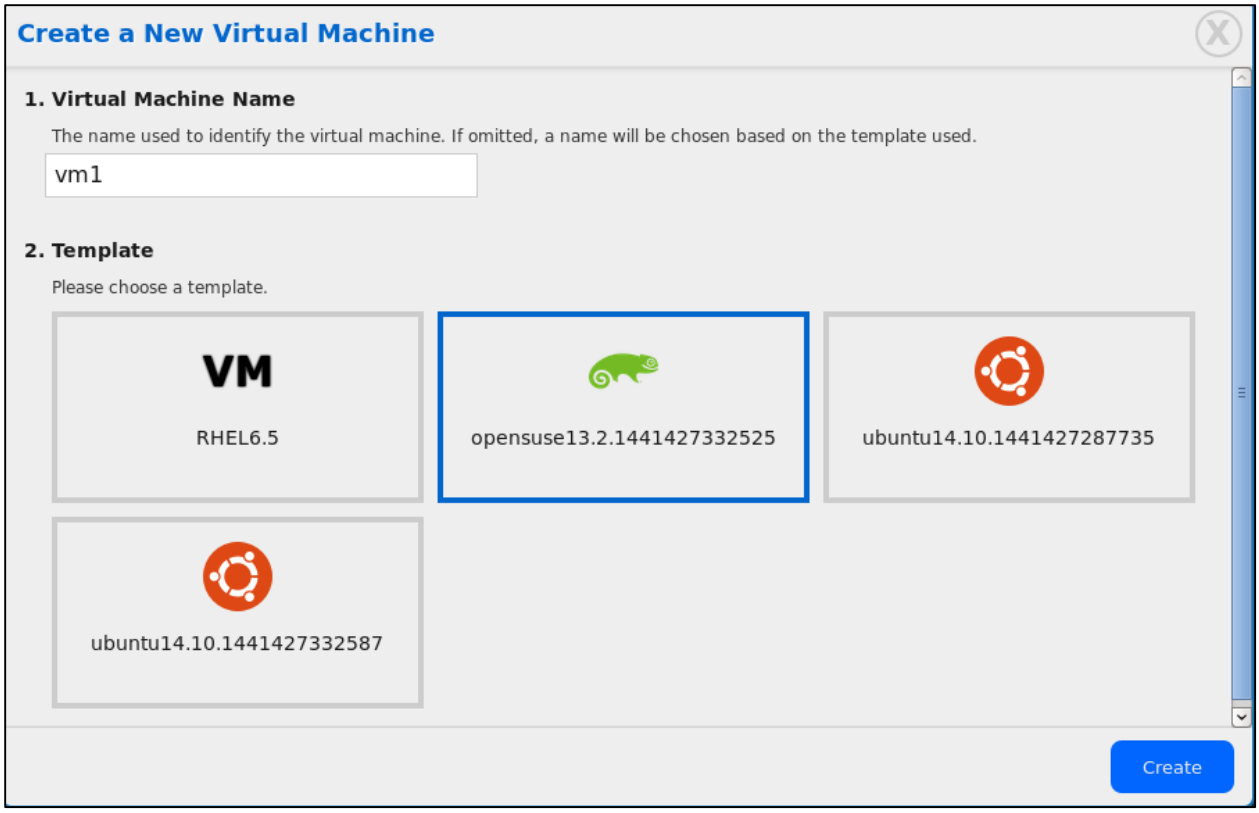
Create a New Virtual Machine
1. Virtual Machine Name
The name used to identify the virtual machine. If omitted, a name will be chosen based on the template used.
vm1
2. Template
Please choose a template.
VM
RHEL6.5
opensuse13.2.1441427332525
ubuntu14.10.1441427287735
ubuntu14.10.1441427332587
Create

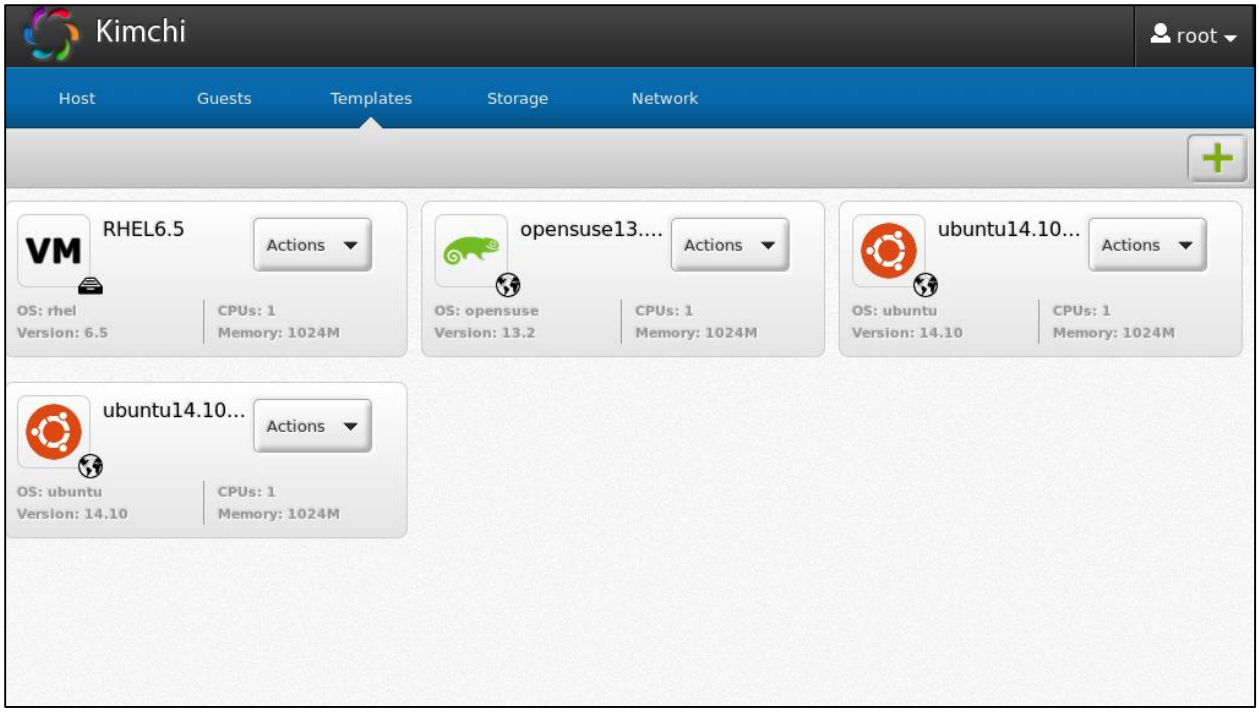
Kimchi
root
Host
Guests
Templates
Storage
Network
RHEL6.5
Actions
OS: rhel
Version: 6.5
CPUs: 1
Memory: 1024M
opensuse13....
Actions
OS: opensuse
Version: 13.2
CPUs: 1
Memory: 1024M
ubuntu14.10...
Actions
OS: ubuntu
Version: 14.10
CPUs: 1
Memory: 1024M
ubuntu14.10...
Actions
OS: ubuntu
Version: 14.10
CPUs: 1
Memory: 1024M

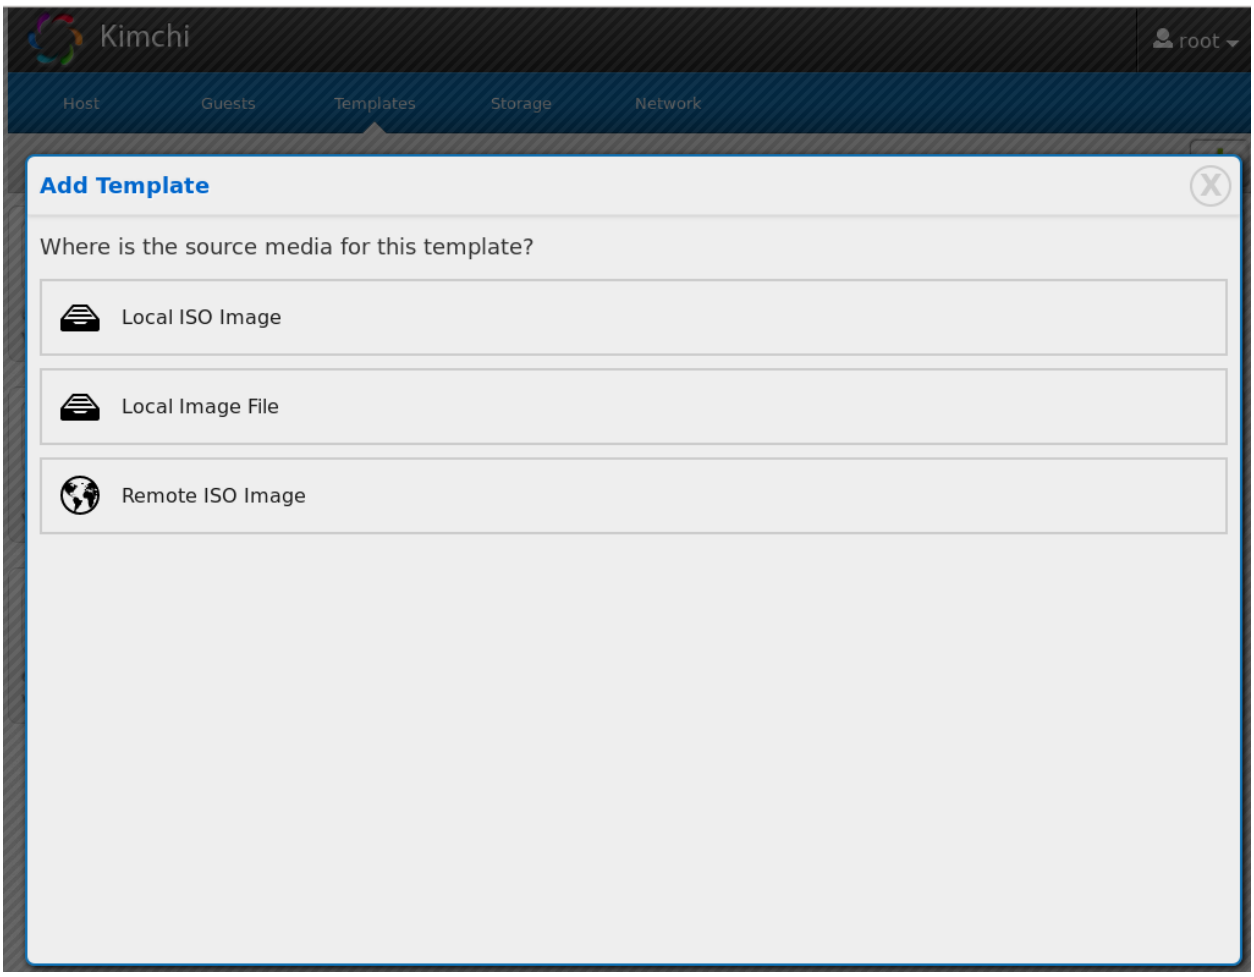
Kimchi
root
Host
Guests
Templates
Storage
Network
Add Template
Where is the source media for this template?
Local ISO Image
Local Image File
Remote ISO Image


Add Template
Local ISO Image
Search ISOs
I want to use a specific ISO file

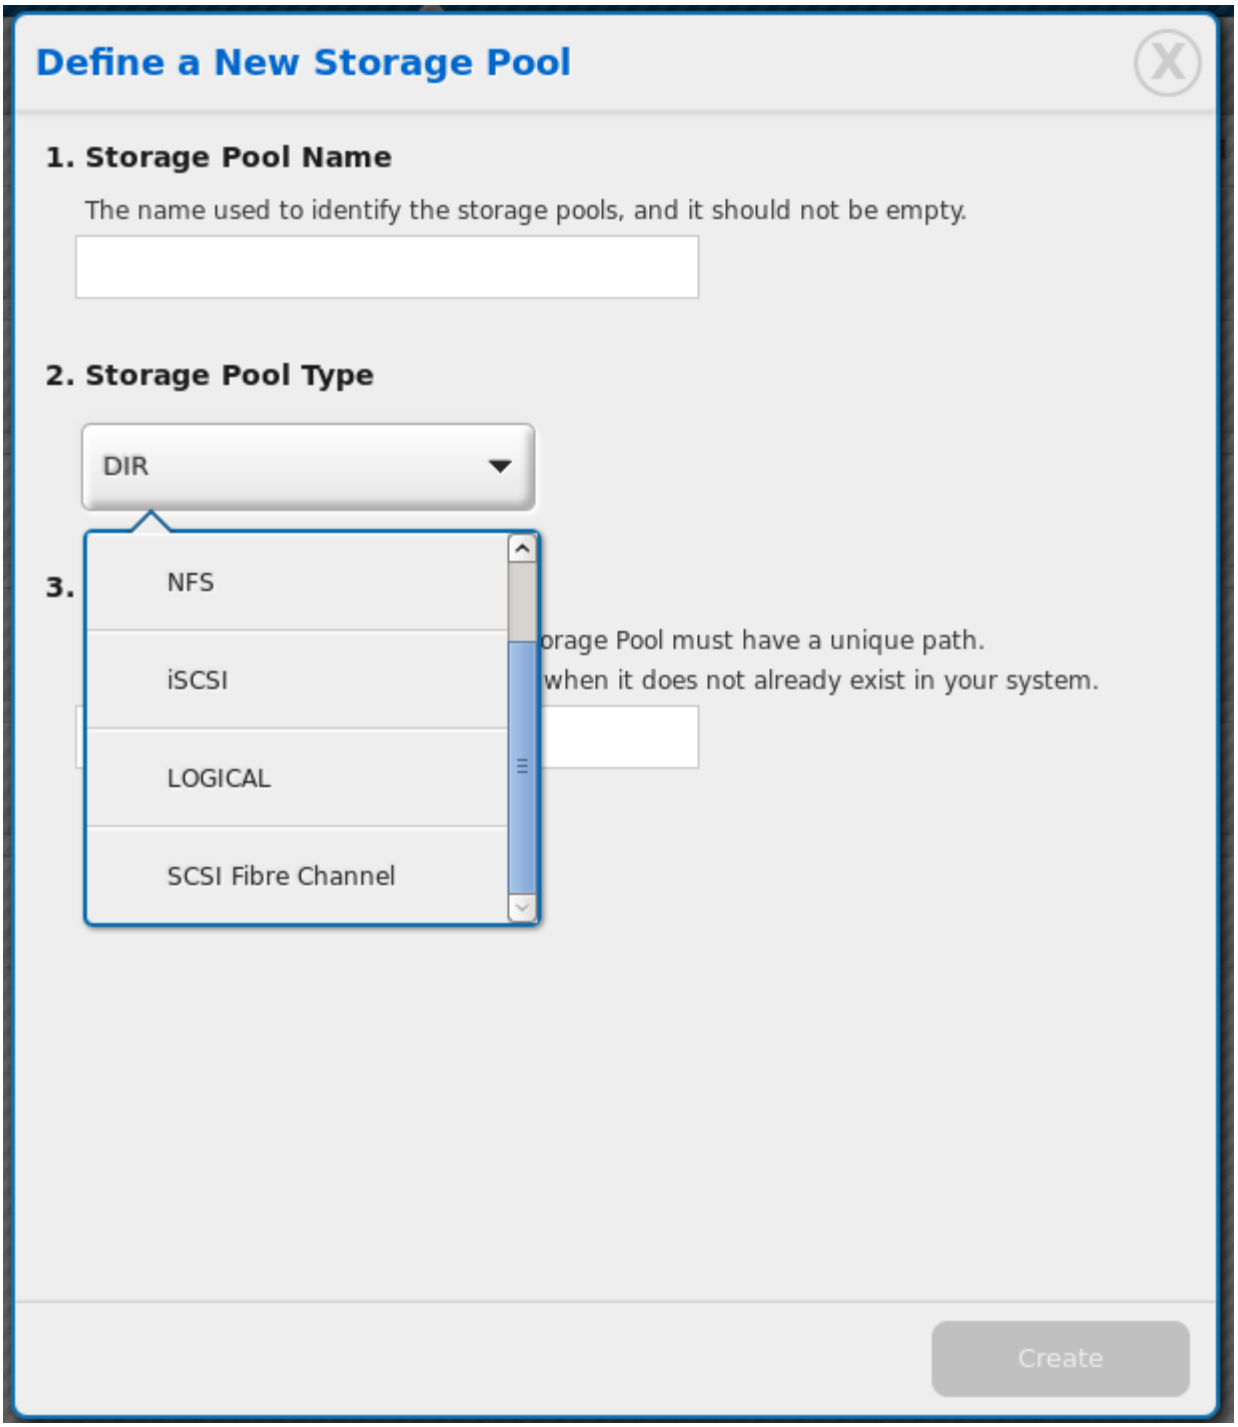
Define a New Storage Pool
1. Storage Pool Name
The name used to identify the storage pools, and it should not be empty.
2. Storage Pool Type
DIR
NFS
iSCSI
LOGICAL
SCSI Fibre Channel
3.
orage Pool must have a unique path.
when it does not already exist in your system.
Create

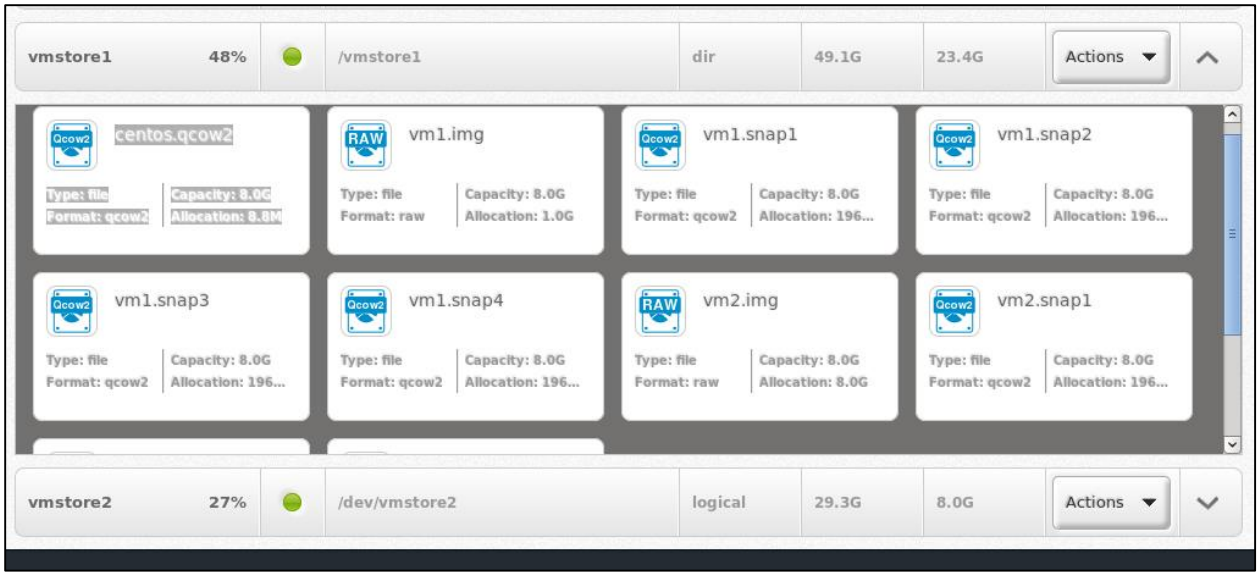
vmstore1 48% /vmstore1 dir 49.1G 23.4G Actions
centos.qcow2
Type: file Format: qcow2 Capacity: 8.0G Allocation: 8.8M
vm1.img
Type: file Format: raw Capacity: 8.0G Allocation: 1.0G
vm1.snap1
Type: file Format: qcow2 Capacity: 8.0G Allocation: 196...
vm1.snap2
Type: file Format: qcow2 Capacity: 8.0G Allocation: 196...
vm1.snap3
Type: file Format: qcow2 Capacity: 8.0G Allocation: 196...
vm1.snap4
Type: file Format: qcow2 Capacity: 8.0G Allocation: 196...
vm2.img
Type: file Format: raw Capacity: 8.0G Allocation: 8.0G
vm2.snap1
Type: file Format: qcow2 Capacity: 8.0G Allocation: 196...
vmstore2 27% /dev/vmstore2 logical 29.3G 8.0G Actions

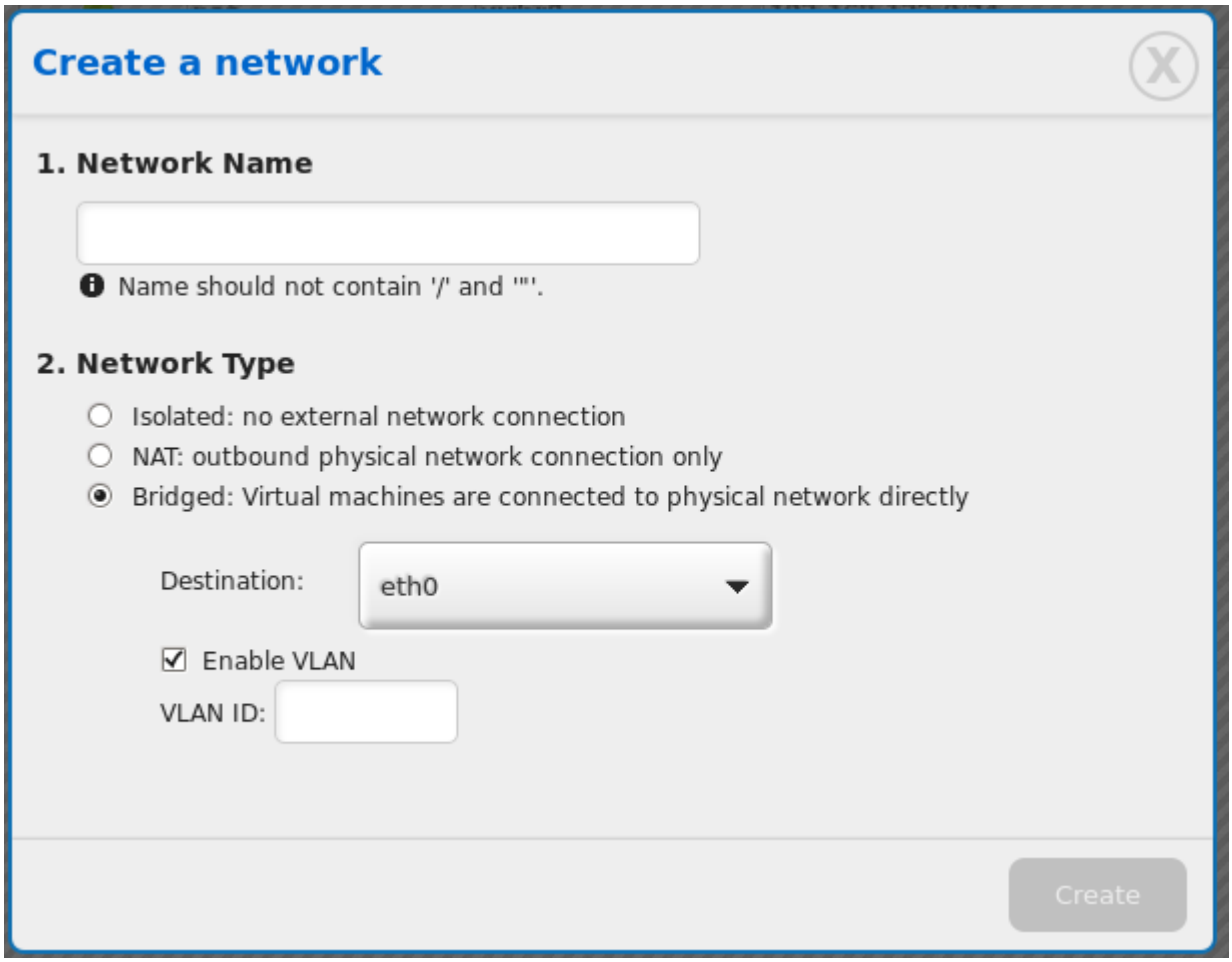
Create a network
1. Network Name
Name should not contain '/' and '"'.
2. Network Type
Isolated: no external network connection
NAT: outbound physical network connection only
Bridged: Virtual machines are connected to physical network directly
Destination:
eth0
Enable VLAN
VLAN ID:
Create

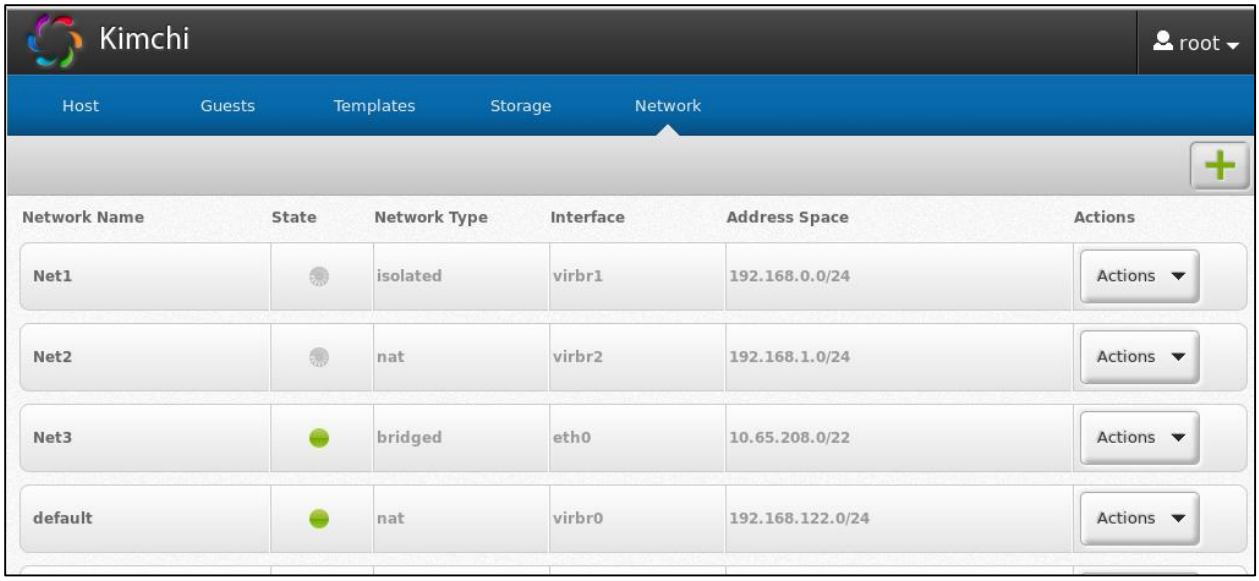
Kimchi
root
Host
Guests
Templates
Storage
Network
Network Name
State
Network Type
Interface
Address Space
Actions
Net1
isolated
virbr1
192.168.0.0/24
Actions
Net2
nat
virbr2
192.168.1.0/24
Actions
Net3
bridged
eth0
10.65.208.0/22
Actions
default
nat
virbr0
192.168.122.0/24
Actions

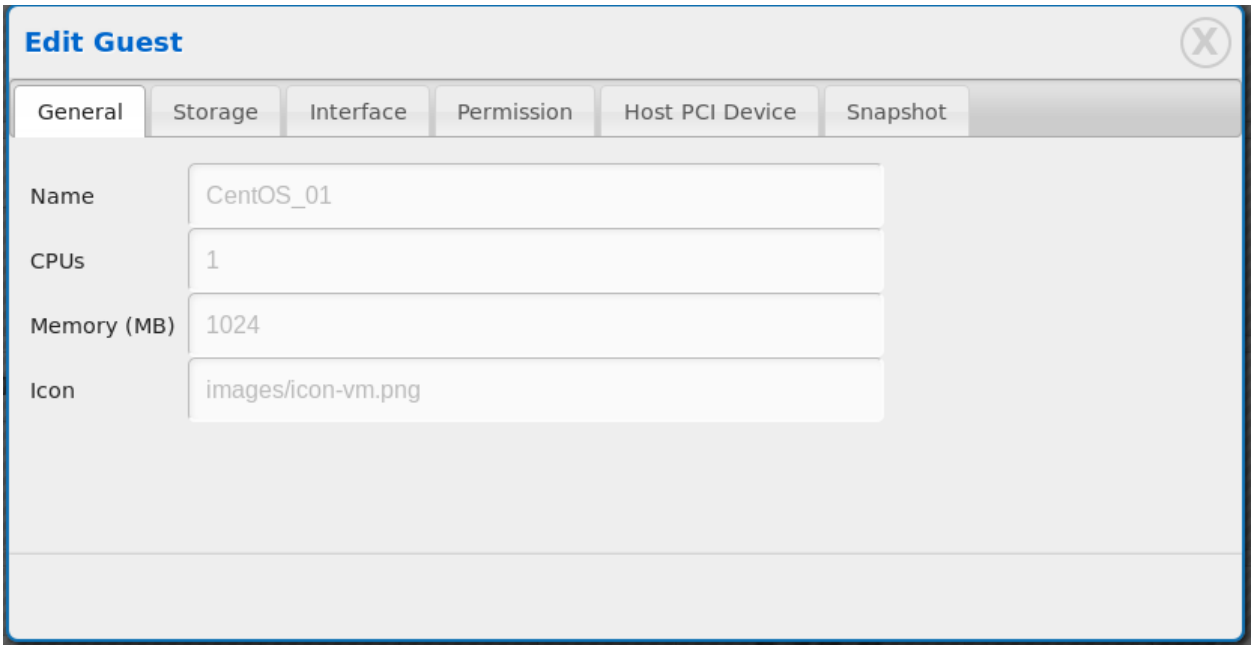
Edit Guest
General
Storage
Interface
Permission
Host PCI Device
Snapshot
Name
CentOS_01
CPUs
1
Memory (MB)
1024
Icon
images/icon-vm.png

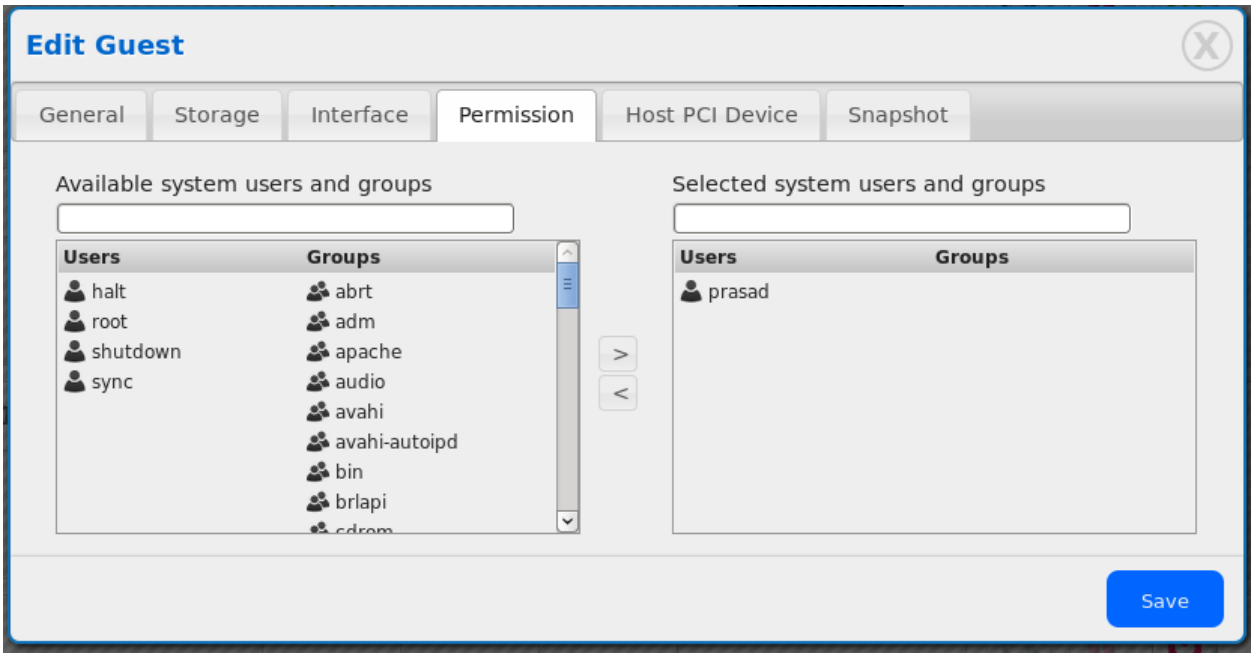
Edit Guest
General
Storage
Interface
Permission
Host PCI Device
Snapshot
Available system users and groups
Users
Groups
halt
root
shutdown
sync
abrt
adm
apache
audio
avahi
avahi-autoipd
bin
brlapi
>
<
Selected system users and groups
Users
Groups
prasad
Save

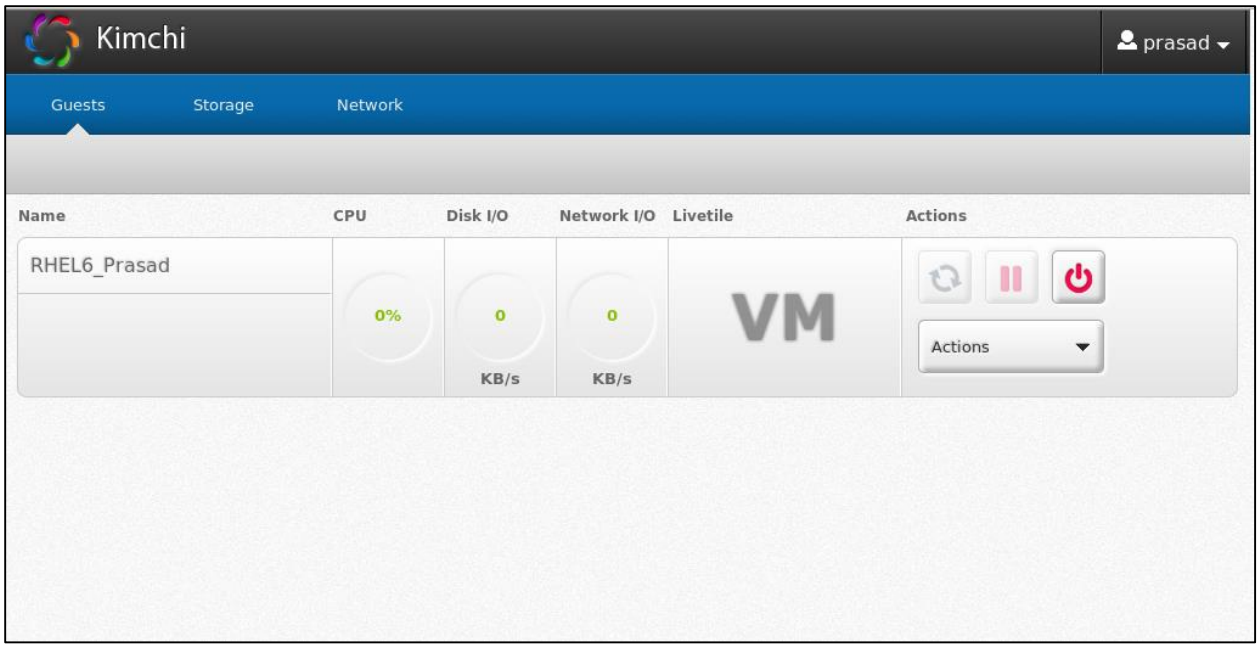
Kimchi
prasad
Guests
Storage
Network
Name
CPU
Disk I/O
Network I/O
Livetile
Actions
RHEL6_Prasad
0%
0
KB/s
0
KB/s
VM
Actions

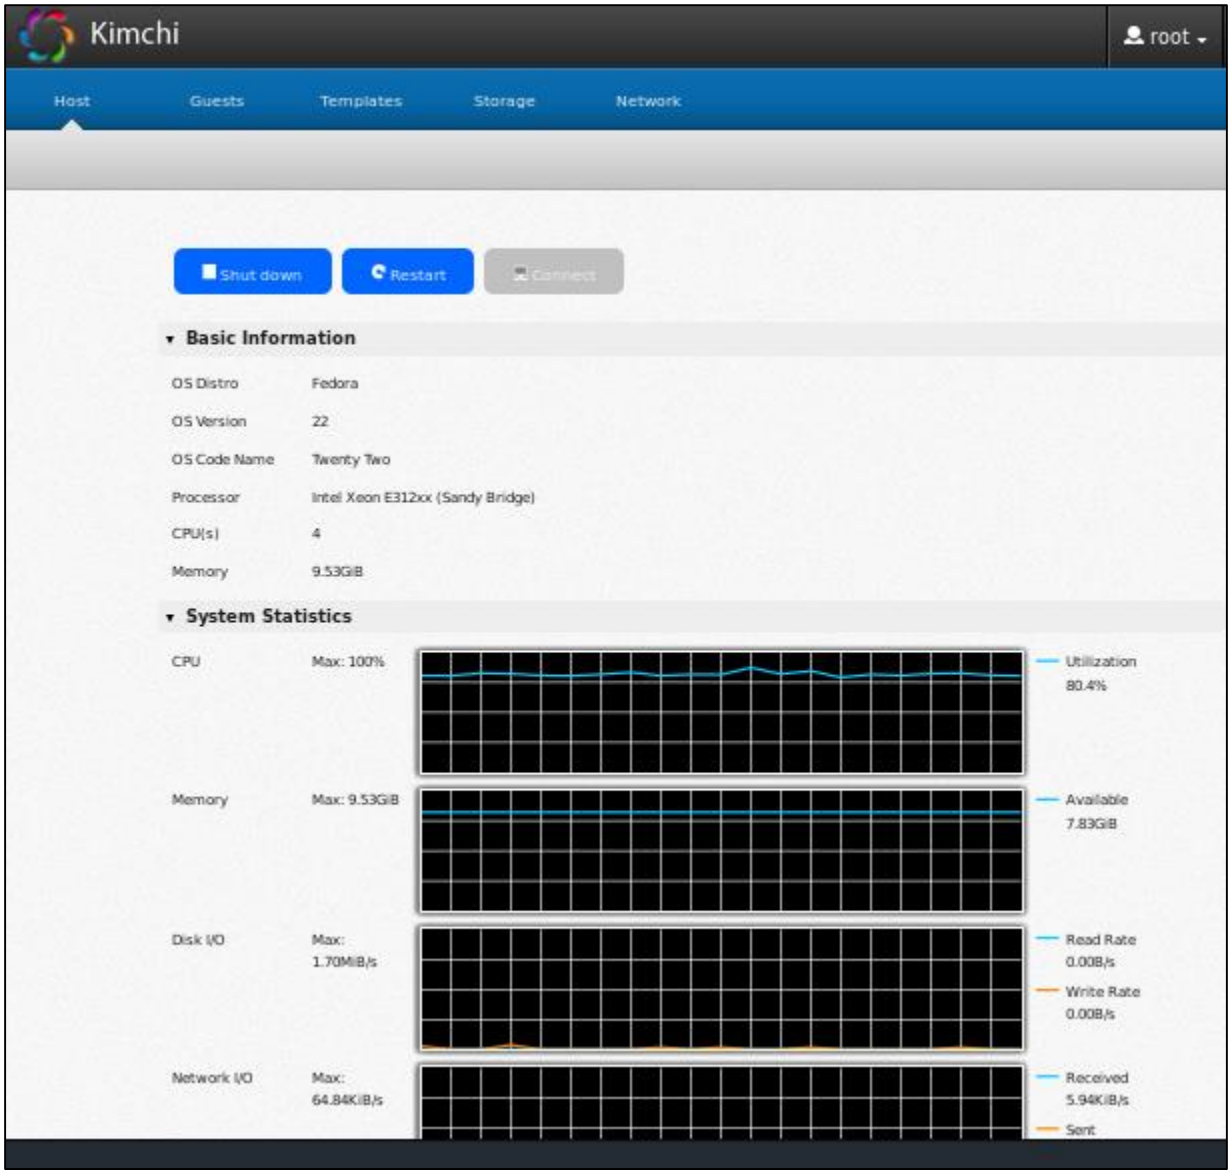
Kimchi
root
Host
Guests
Templates
Storage
Network
Shut down
Restart
Connect
Basic Information
OS Distro
Fedora
OS Version
22
OS Code Name
Twenty Two
Processor
Intel Xeon E312xx (Sandy Bridge)
CPU(s)
4
Memory
9.53GiB
System Statistics
CPU
Max: 100%
Utilization
80.4%
Memory
Max: 9.53GiB
Available
7.83GiB
Disk I/O
Max:
1.70MiB/s
Read Rate
0.00B/s
Write Rate
0.00B/s
Network I/O
Max:
64.84KiB/s
Received
5.94KiB/s
Sent

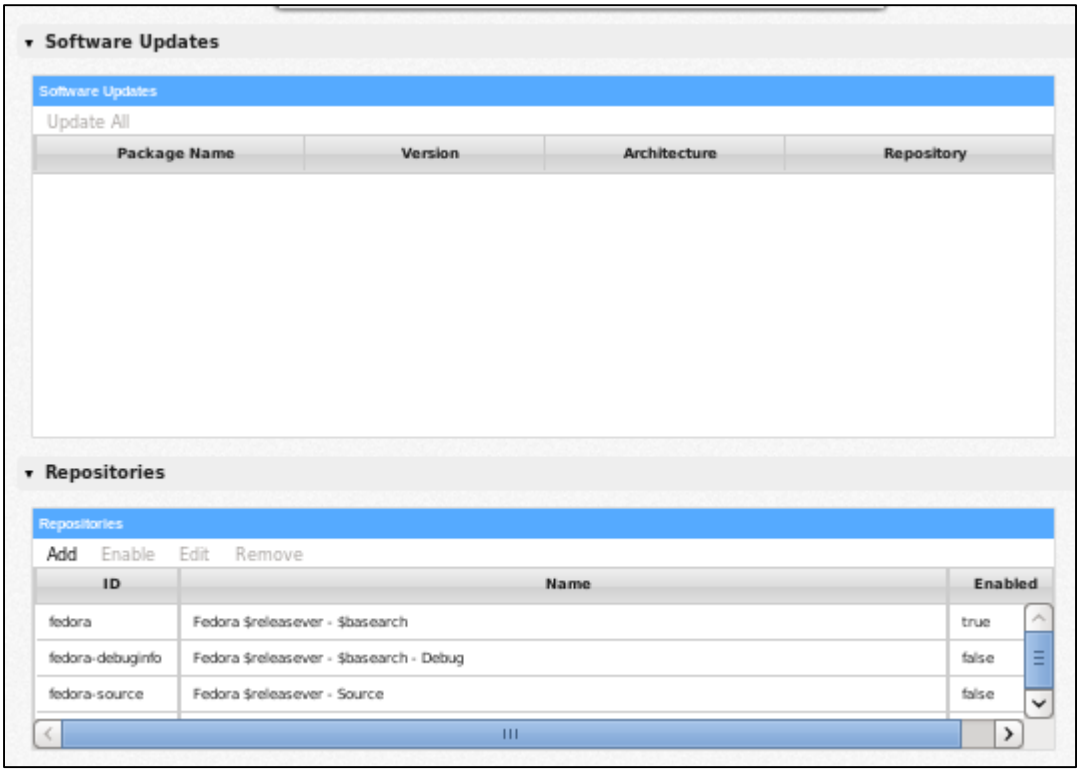
Software Updates
Software Updates
Update All
Package Name
Version
Architecture
Repository
Repositories
Repositories
Add
Enable
Edit
Remove
ID
Name
Enabled
fedora
Fedora $releasever - $basearch
true
fedora-debuginfo
Fedora $releasever - $basearch - Debug
false
fedora-source
Fedora $releasever - Source
false

Chapter 9: Software-Defied Networking for KVM Virtualization

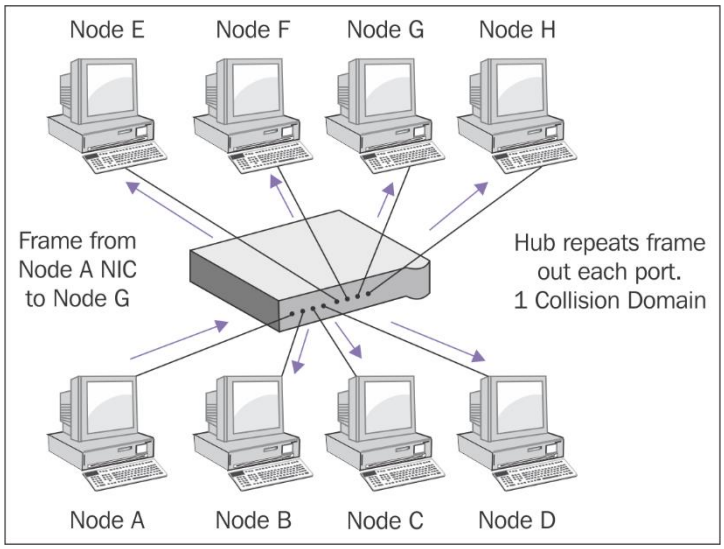

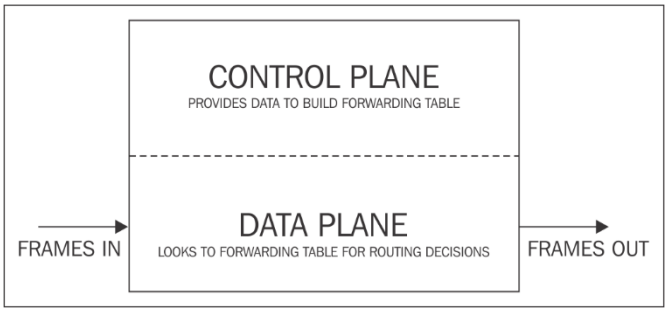

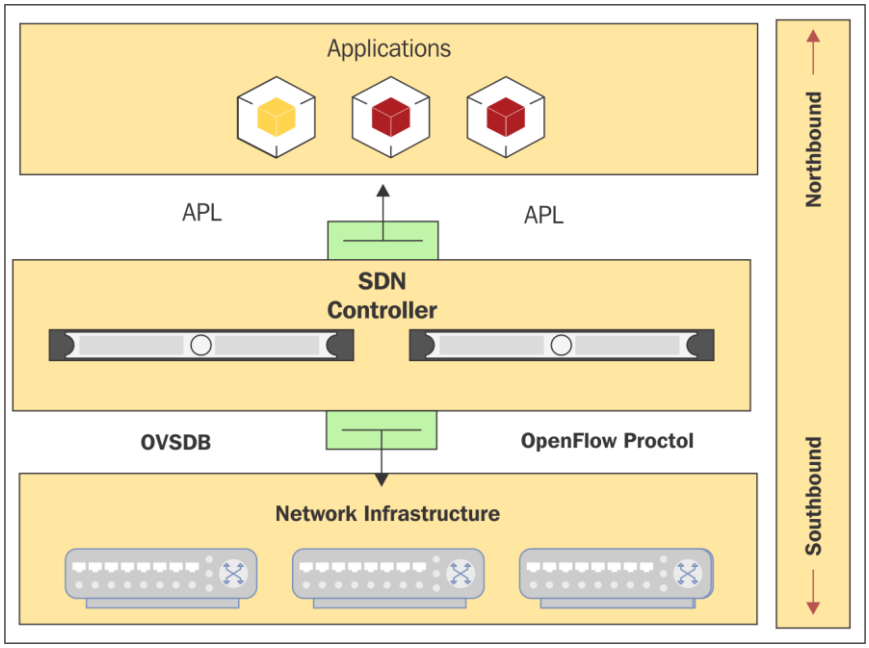

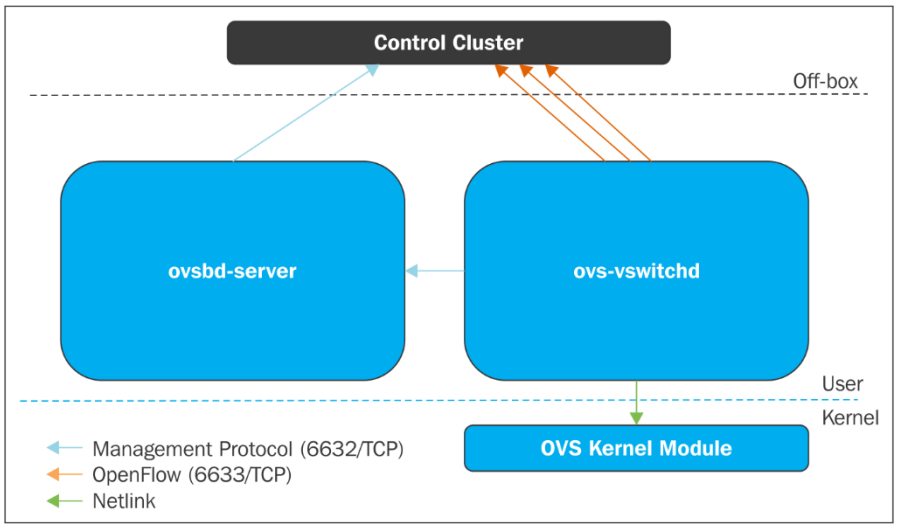
Control Cluster
Off-box
ovsbd-server
ovs-vswitchd
User
Kernel
OVS Kernel Module
Management Protocol (6632/TCP)
OpenFlow (6633/TCP)
Netlink

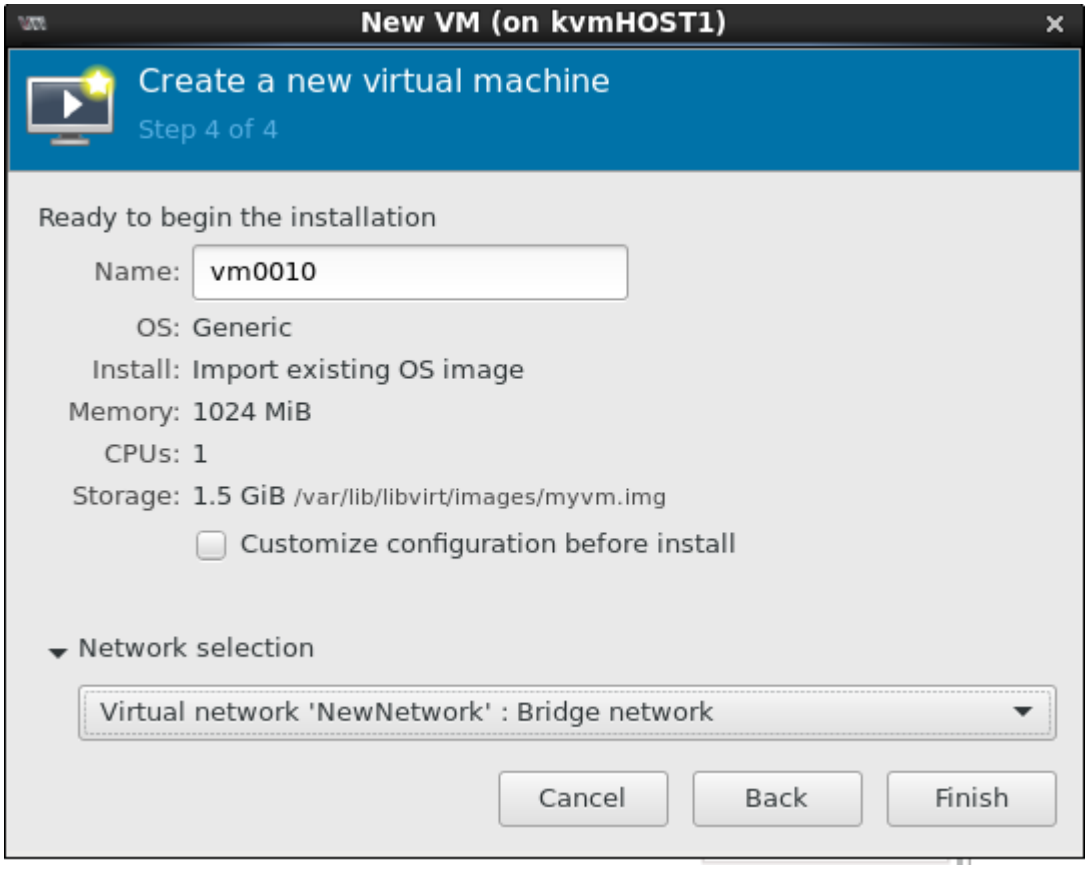
New VM (on kvmHOST1)
Create a new virtual machine
Step 4 of 4
Ready to begin the installation
Name: vm0010
OS: Generic
Install: Import existing OS image
Memory: 1024 MiB
CPUs: 1
Storage: 1.5 GiB /var/lib/libvirt/images/myvm.img
Customize configuration before install
Network selection
Virtual network 'NewNetwork' : Bridge network
Cancel
Back
Finish

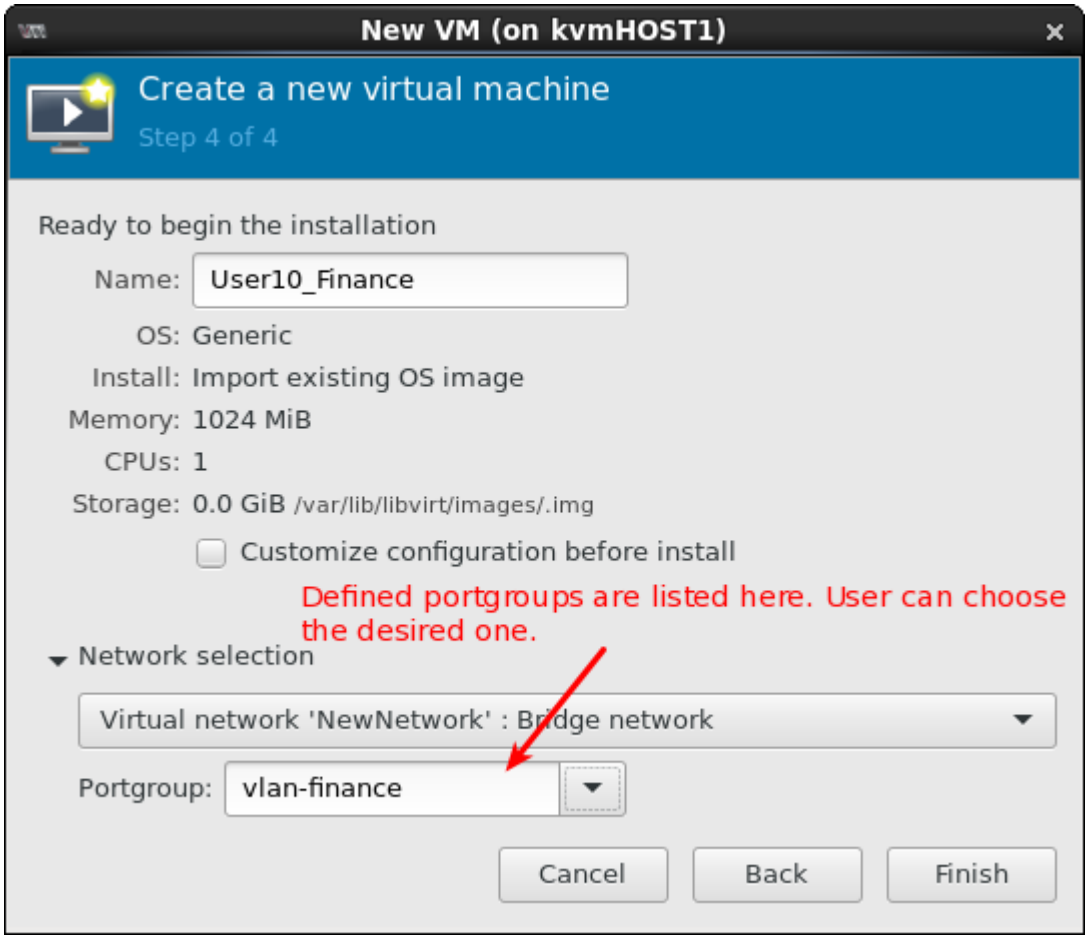
New VM (on kvmHOST1)
Create a new virtual machine
Step 4 of 4
Ready to begin the installation
Name: User10_Finance
OS: Generic
Install: Import existing OS image
Memory: 1024 MiB
CPUs: 1
Storage: 0.0 GiB /var/lib/libvirt/images/.img
Customize configuration before install
Defined portgroups are listed here. User can choose the desired one.
Network selection
Virtual network 'NewNetwork' : Bridge network
Portgroup: vlan-finance
Cancel
Back
Finish

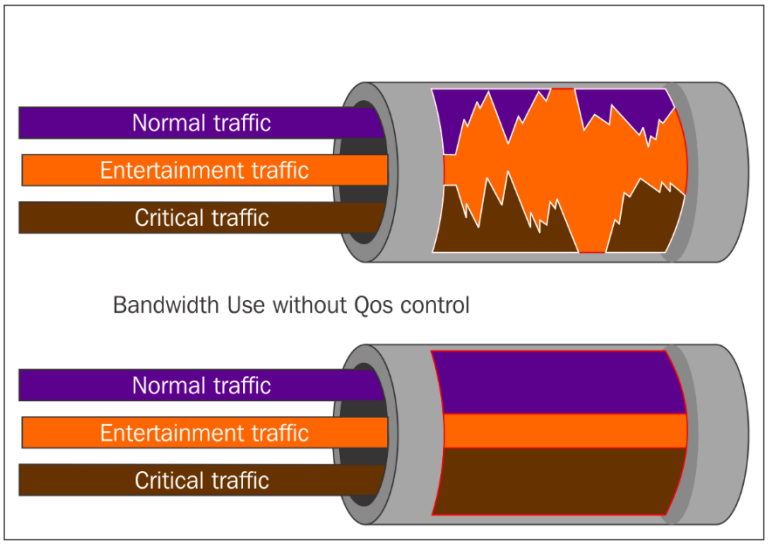
Normal traffic
Entertainment traffic
Critical traffic
Bandwidth Use without Qos control
Normal traffic
Entertainment traffic
Critical traffic

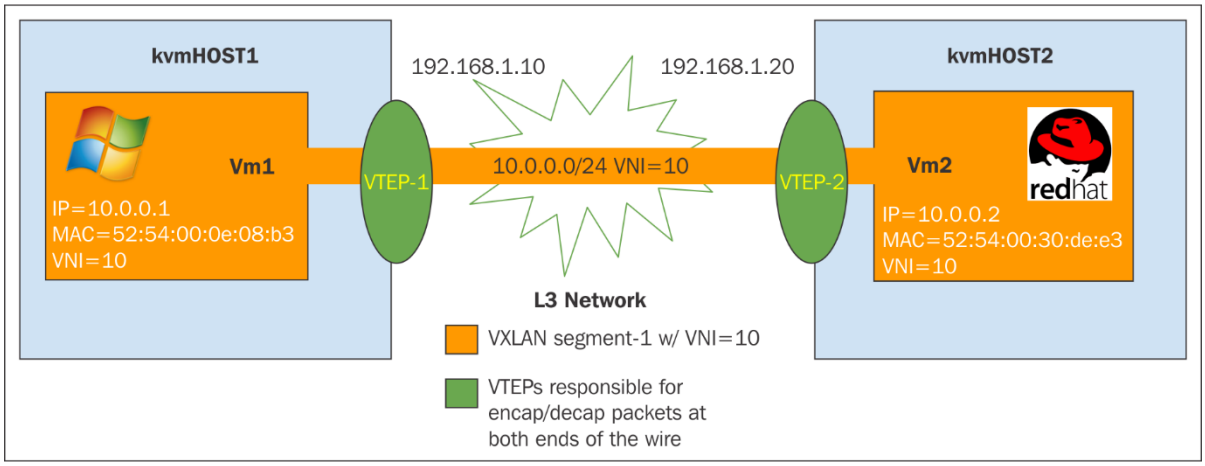
kvmHOST1
192.168.1.10
192.168.1.20
kvmHOST2
Vm1
IP=10.0.0.1
MAC=52:54:00:0e:08:b3
VNI=10
VTEP-1
10.0.0.0/24 VNI=10
VTEP-2
Vm2
IP=10.0.0.2
MAC=52:54:00:30:de:e3
VNI=10
L3 Network
VXLAN segment-1 w/ VNI=10
VTEPs responsible for encap/decap packets at both ends of the wire

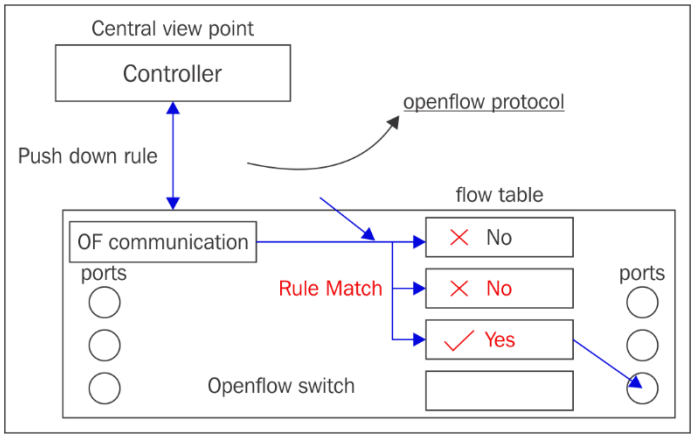
Central view point
Controller
openflow protocol
Push down rule
flow table
OF communication
No
ports
Rule Match
No
ports
Yes
Openflow switch

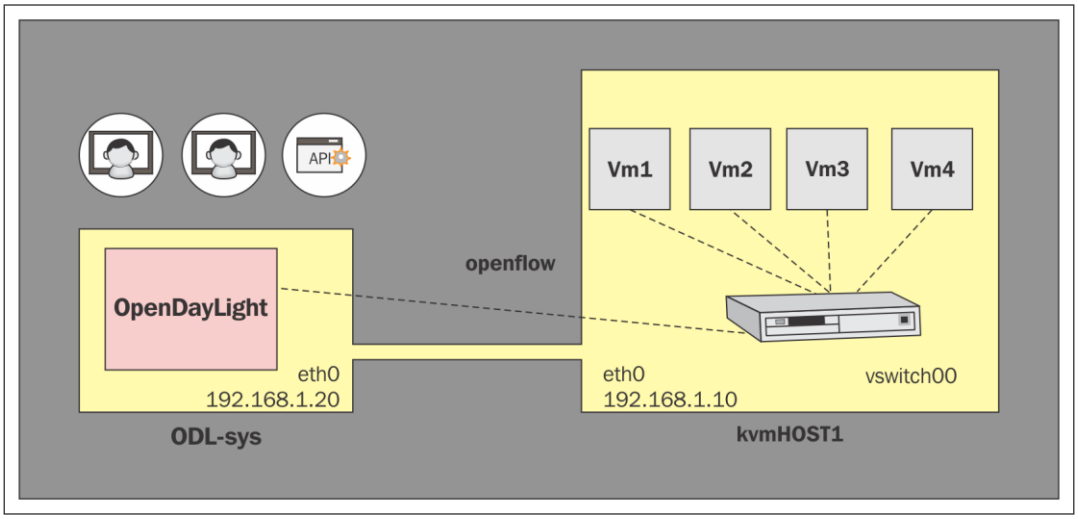
API
Vm1
Vm2
Vm3
Vm4
OpenDayLight
openflow
eth0
192.168.1.20
ODL-sys
eth0
192.168.1.10
vswitch00
kvmHOST1

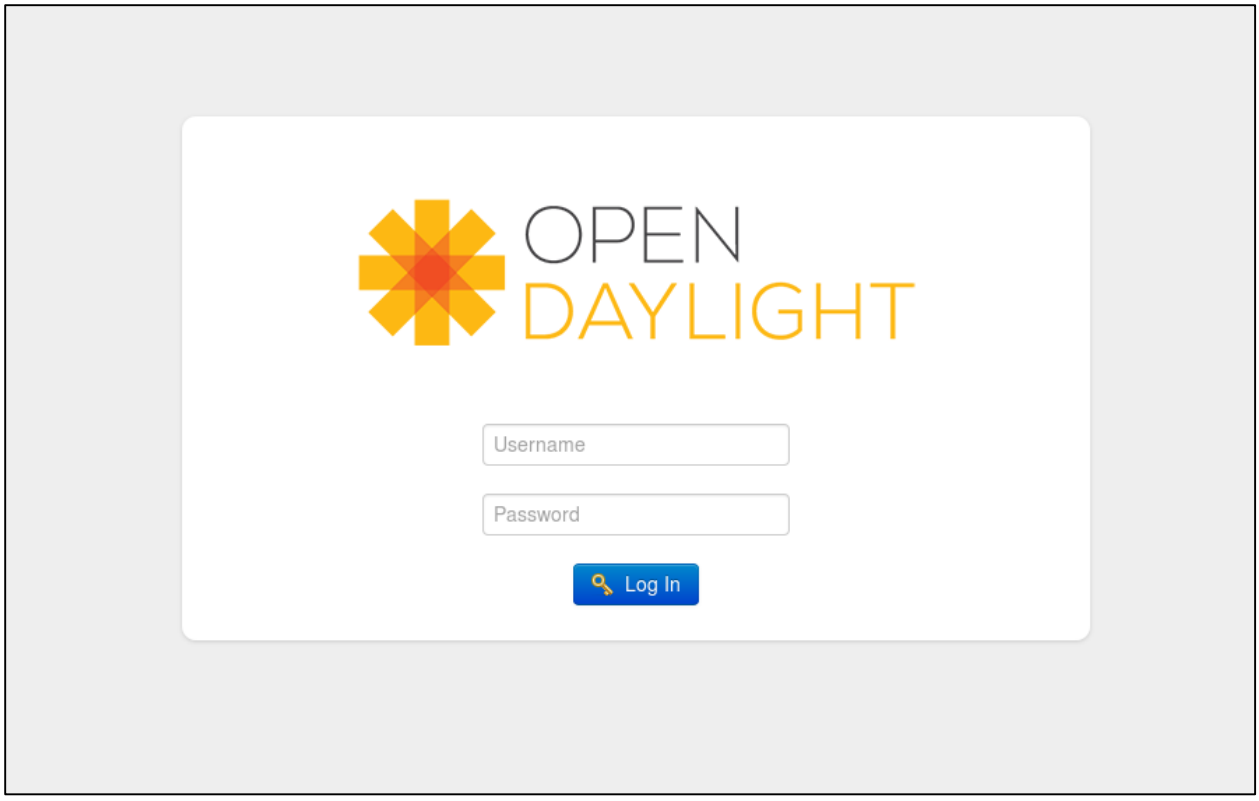
OPEN
DAYLIGHT
Username
Password
Log In

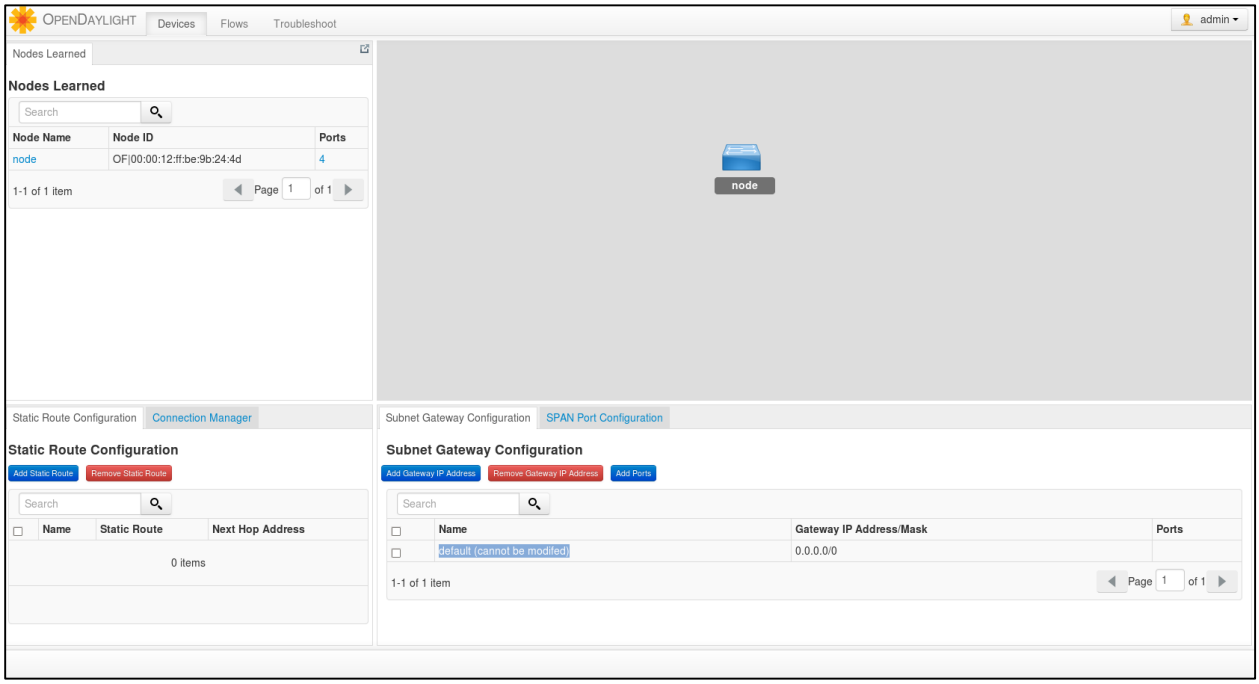
OPENDAYLIGHT
Devices
Flows
Troubleshoot
admin
Nodes Learned
Nodes Learned
Search
Node Name
Node ID
Ports
node
OF|00:00:12:ff:be:9b:24:4d
4
1-1 of 1 item
Page 1 of 1
node
Static Route Configuration
Connection Manager
Static Route Configuration
Add Static Route
Remove Static Route
Search
Name
Static Route
Next Hop Address
0 items
Subnet Gateway Configuration
SPAN Port Configuration
Subnet Gateway Configuration
Add Gateway IP Address
Remove Gateway IP Address
Add Ports
Search
Name
Gateway IP Address/Mask
Ports
default (cannot be modified)
0.0.0.0/0
1-1 of 1 item
Page 1 of 1

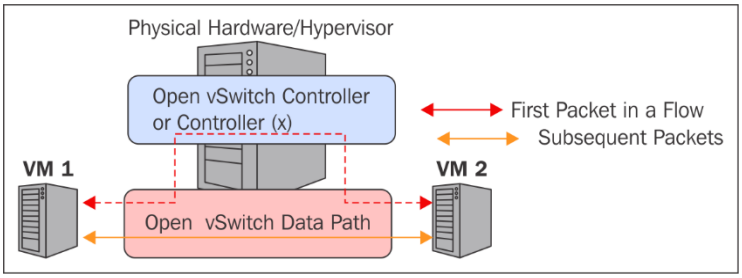
Physical Hardware/Hypervisor
Open vSwitch Controller
or Controller (x)
First Packet in a Flow
Subsequent Packets
VM 1
VM 2
Open vSwitch Data Path

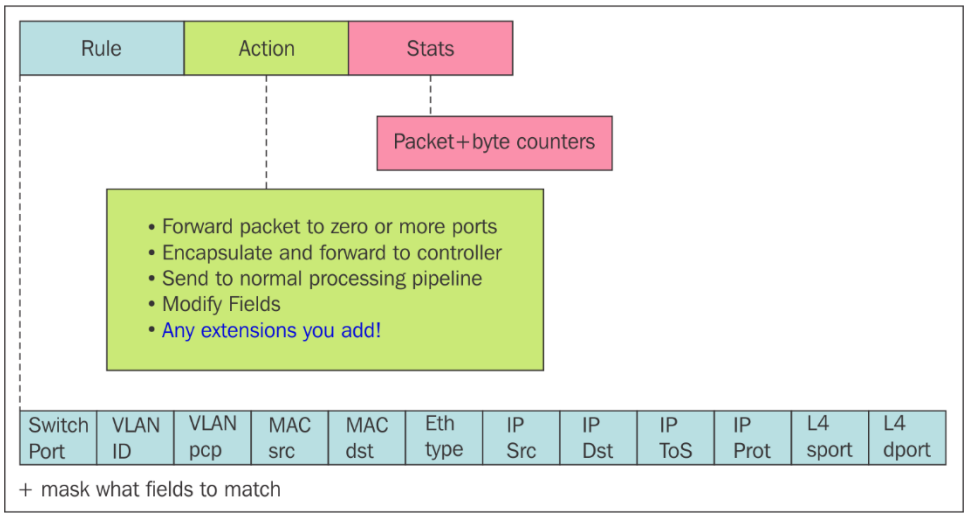
Rule
Action
Stats
Packet+byte counters
• Forward packet to zero or more ports
• Encapsulate and forward to controller
• Send to normal processing pipeline
• Modify Fields
• Any extensions you add!
Switch Port
VLAN ID
VLAN pcp
MAC src
MAC dst
Eth type
IP Src
IP Dst
IP ToS
IP Prot
L4 sport
L4 dport
+ mask what fields to match

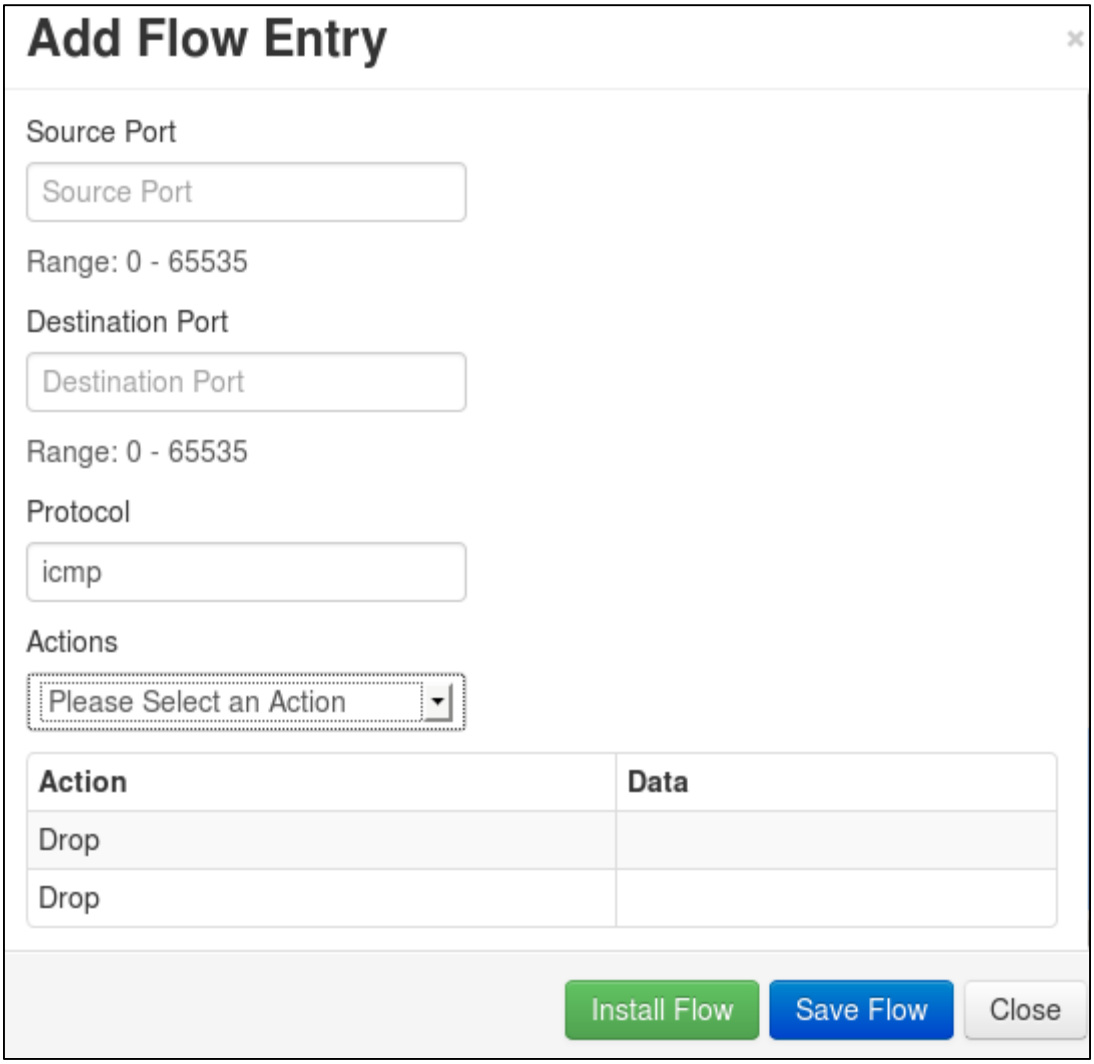
Add Flow Entry
Source Port
Source Port
Range: 0 - 65535
Destination Port
Destination Port
Range: 0 - 65535
Protocol
icmp
Actions
Please Select an Action
Action
Data
Drop
Drop
Install Flow
Save Flow
Close

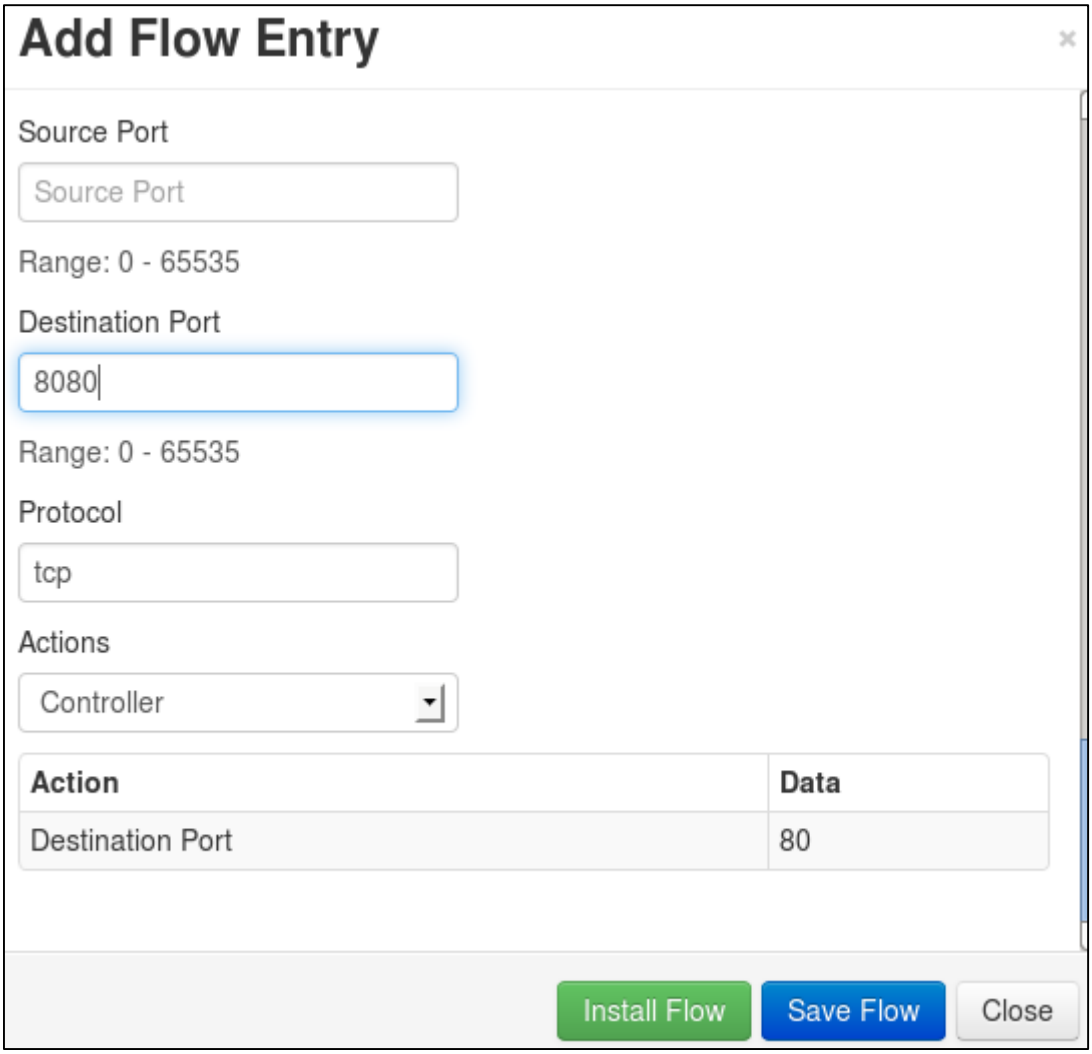
Add Flow Entry
Source Port
Source Port
Range: 0 - 65535
Destination Port
8080
Range: 0 - 65535
Protocol
tcp
Actions
Controller
Action
Data
Destination Port
80
Install Flow
Save Flow
Close

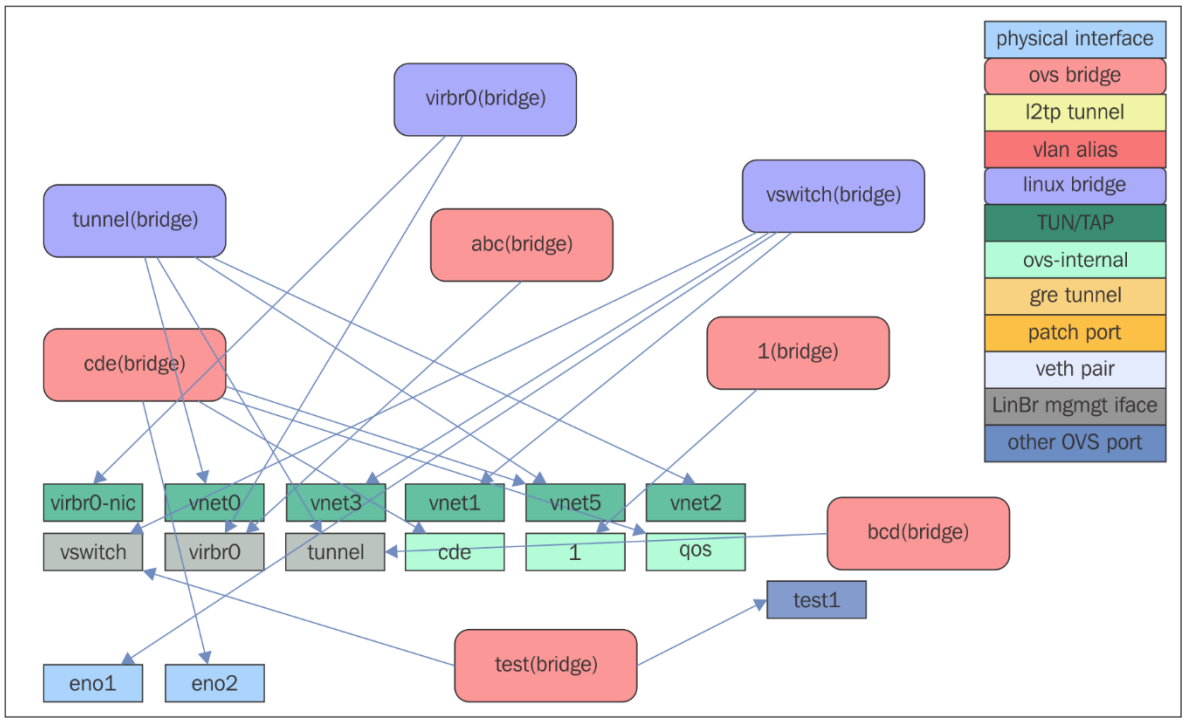
physical interface
ovs bridge
l2tp tunnel
vlan alias
linux bridge
TUN/TAP
ovs-internal
gre tunnel
patch port
veth pair
LinBr mgmgt iface
other OVS port
virbr0(bridge)
tunnel(bridge)
abc(bridge)
vswitch(bridge)
cde(bridge)
1(bridge)
virbr0-nic
vnet0
vnet3
vnet1
vnet5
vnet2
vswitch
virbr0
tunnel
cde
1
qos
bcd(bridge)
test1
test(bridge)
eno1
eno2

Chapter 10: Installing and Configuring the Virtual Datacenter Using oVirt

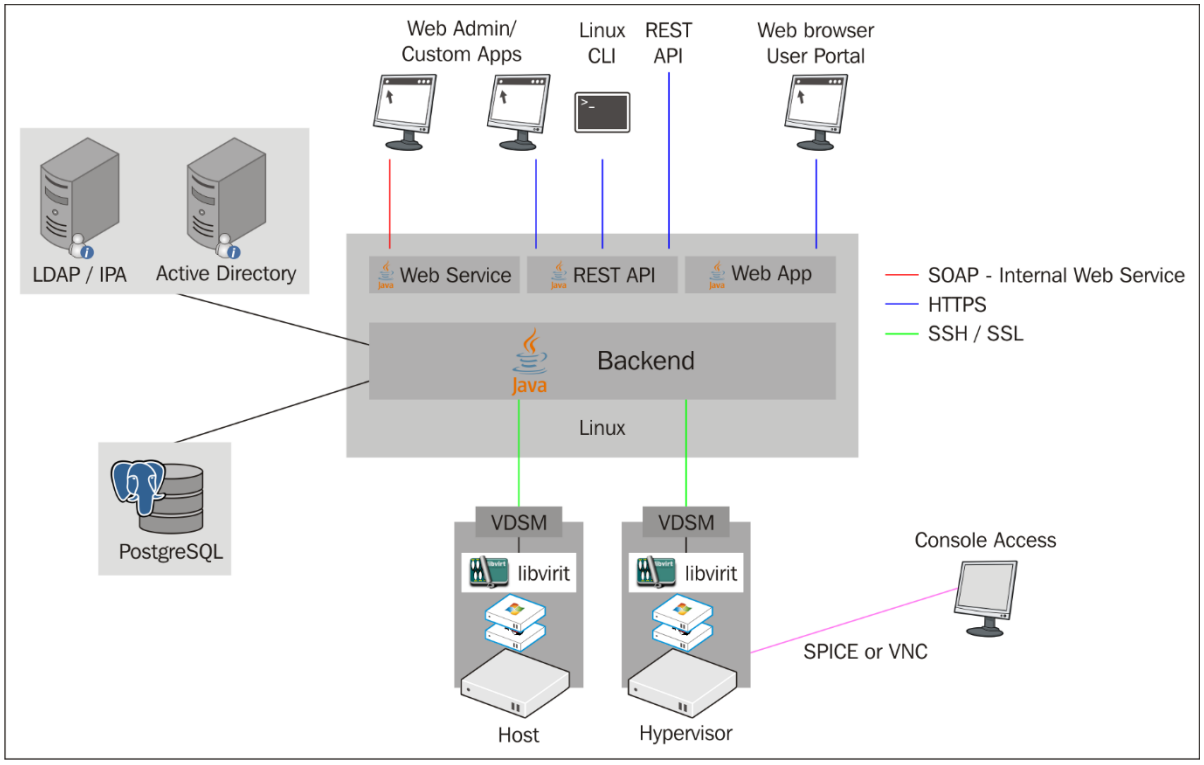

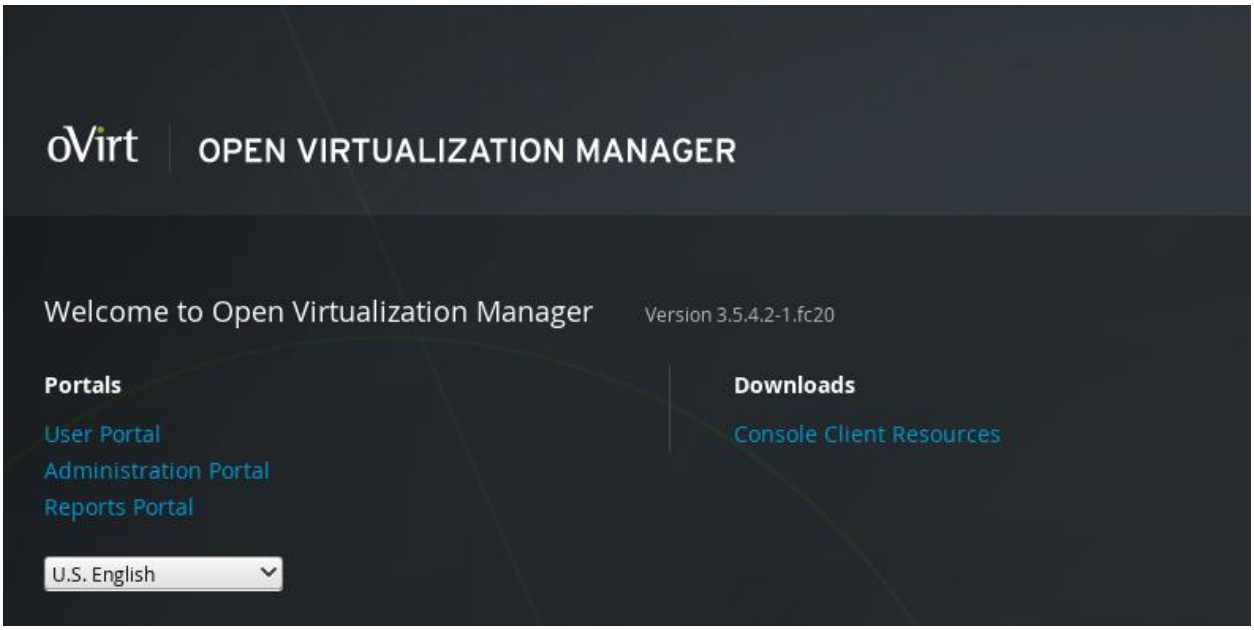

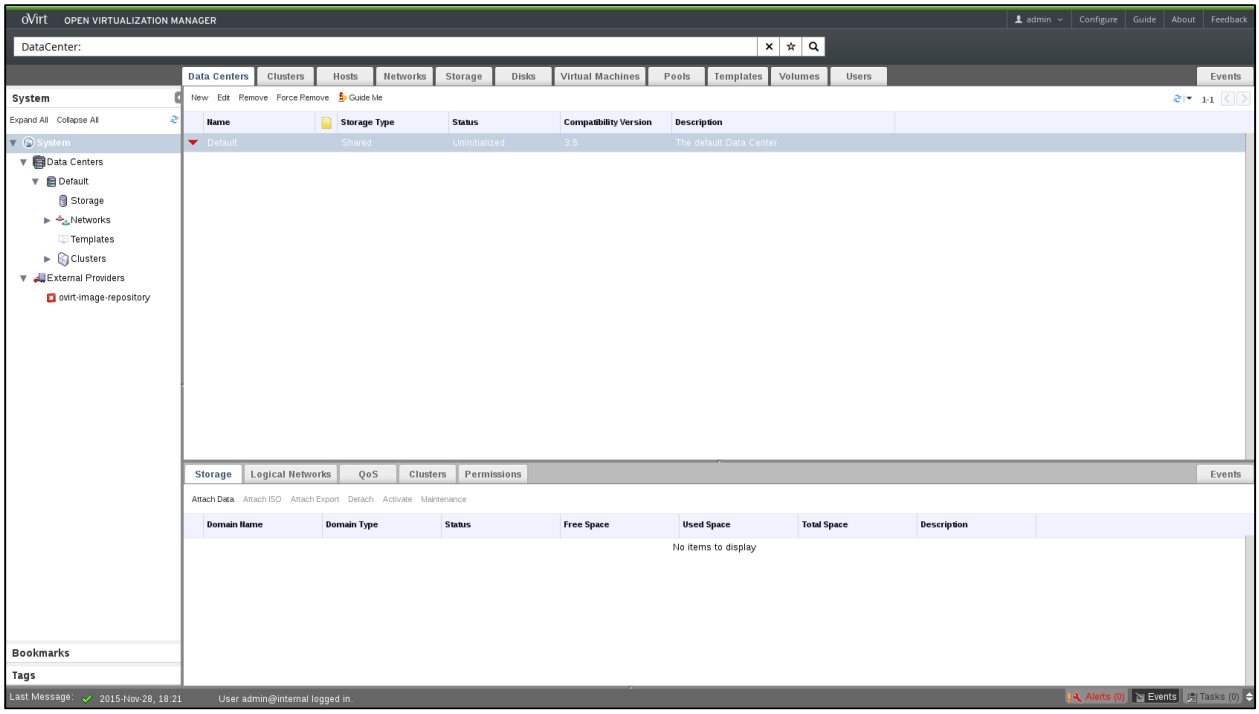
oVirt
OPEN VIRTUALIZATION MANAGER
admin
Configure
Guide
About
Feedback
DataCenter:
Data Centers
Clusters
Hosts
Networks
Storage
Disks
Virtual Machines
Pools
Templates
Volumes
Users
Events
System
Expand All
Collapse All
New
Edit
Remove
Force Remove
Guide Me
Name
Storage Type
Status
Compatibility Version
Description
Default
Shared
Uninitialized
3.5
The default Data Center
Data Centers
Default
Storage
Networks
Templates
Clusters
External Providers
ovirt-image-repository
Storage
Logical Networks
QoS
Clusters
Permissions
Events
Attach Data
Attach ISO
Attach Export
Detach
Activate
Maintenance
Domain Name
Domain Type
Status
Free Space
Used Space
Total Space
Description
No items to display
Bookmarks
Tags
Last Message:
2015-Nov-28, 18:21
User admin@internal logged in.
Alerts (0)
Events
Tasks (0)

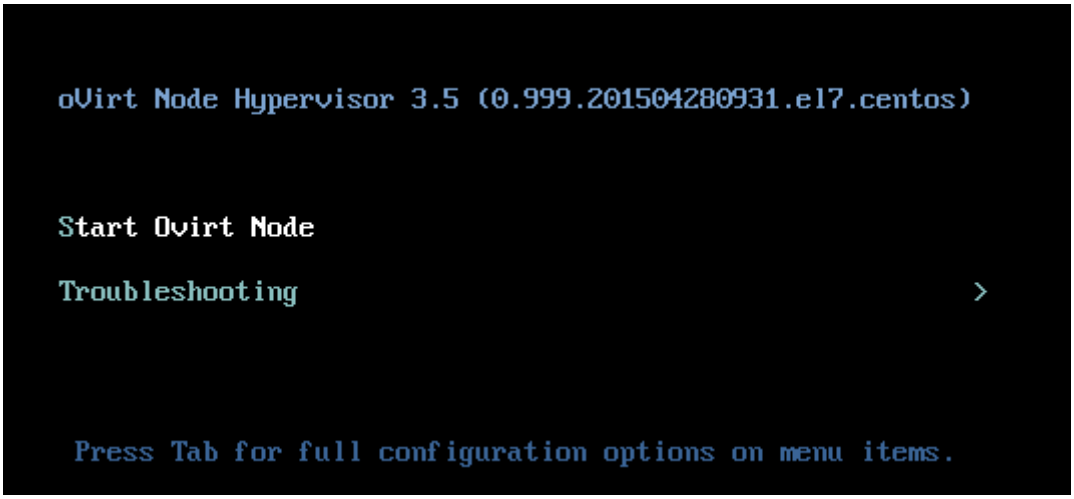
oVirt Node Hypervisor 3.5 (0.999.201504280931.el7.centos)
Start Ovirt Node
Troubleshooting
Press Tab for full configuration options on menu items.

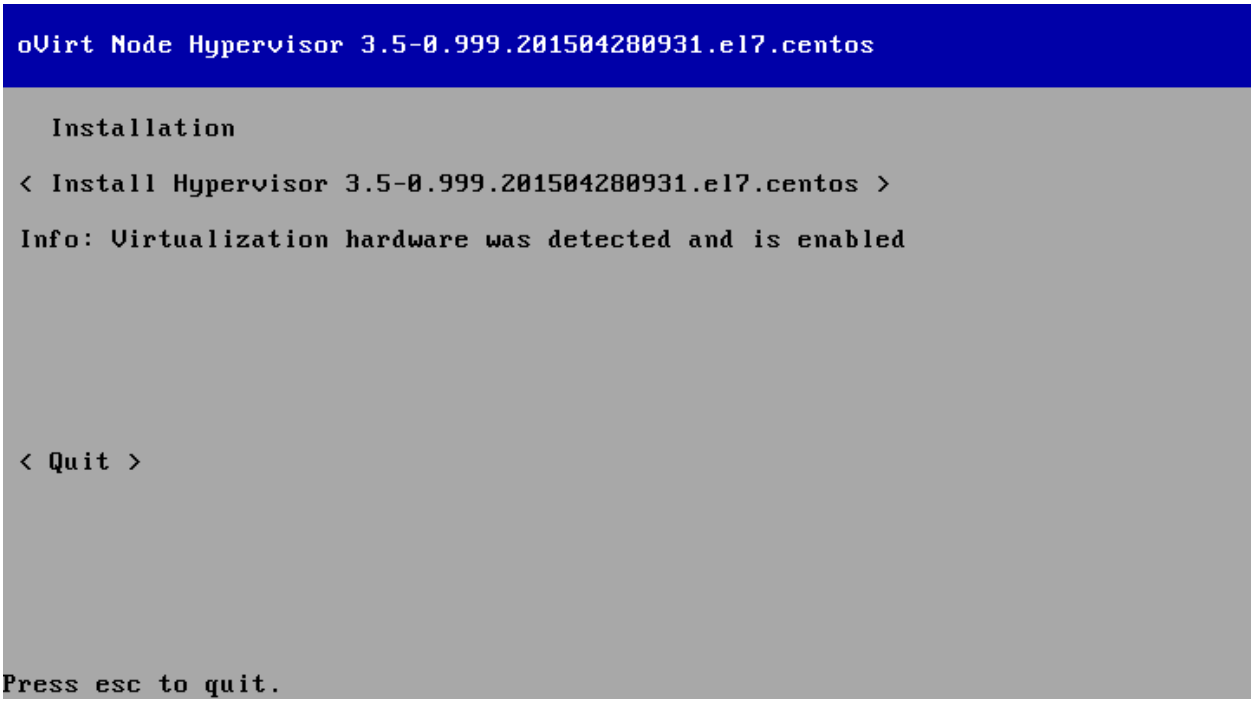
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Installation
< Install Hypervisor 3.5-0.999.201504280931.el7.centos >
Info: Virtualization hardware was detected and is enabled
< Quit >
Press esc to quit.

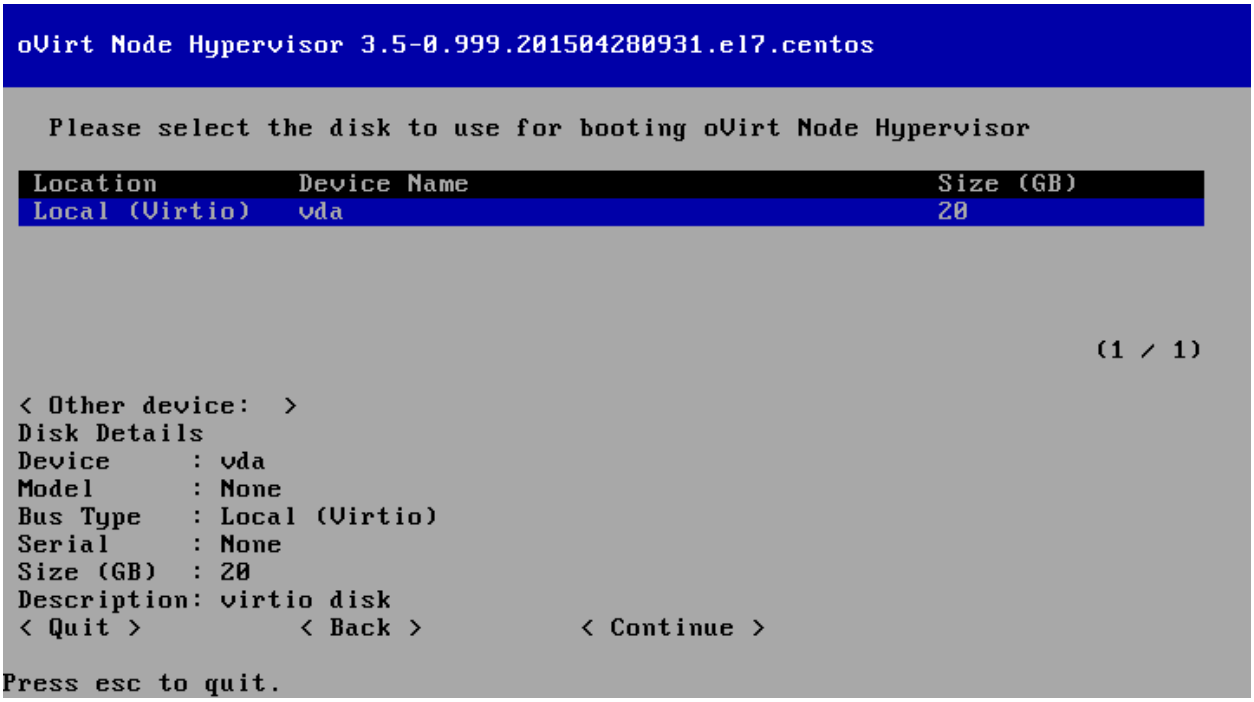
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Please select the disk to use for booting oVirt Node Hypervisor
Location
Device Name
Size (GB)
Local (Virtio)
vda
20
(1 / 1)
< Other device: >
Disk Details
Device : vda
Model : None
Bus Type : Local (Virtio)
Serial : None
Size (GB) : 20
Description: virtio disk
< Quit >
< Back >
< Continue >
Press esc to quit.

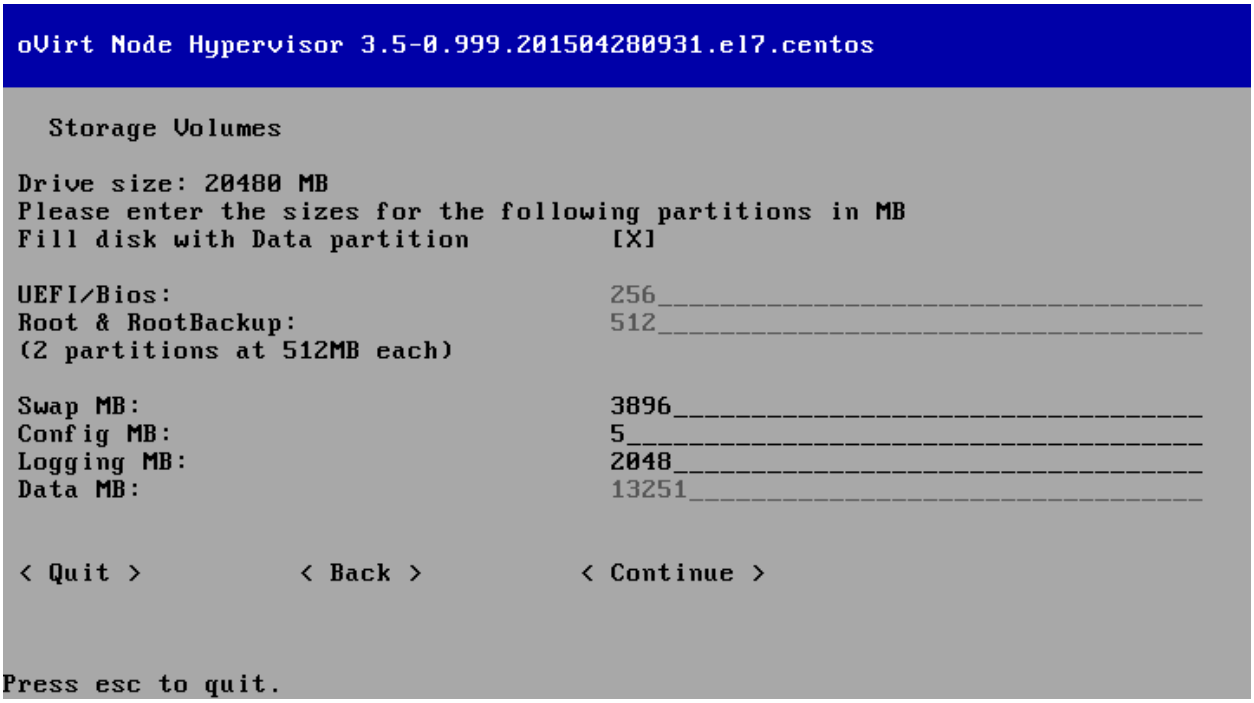
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Storage Volumes
Drive size: 20480 MB
Please enter the sizes for the following partitions in MB
Fill disk with Data partition [X]
UEFI/Bios: 256
Root & RootBackup: 512
(2 partitions at 512MB each)
Swap MB: 3896
Config MB: 5
Logging MB: 2048
Data MB: 13251
< Quit > < Back > < Continue >
Press esc to quit.

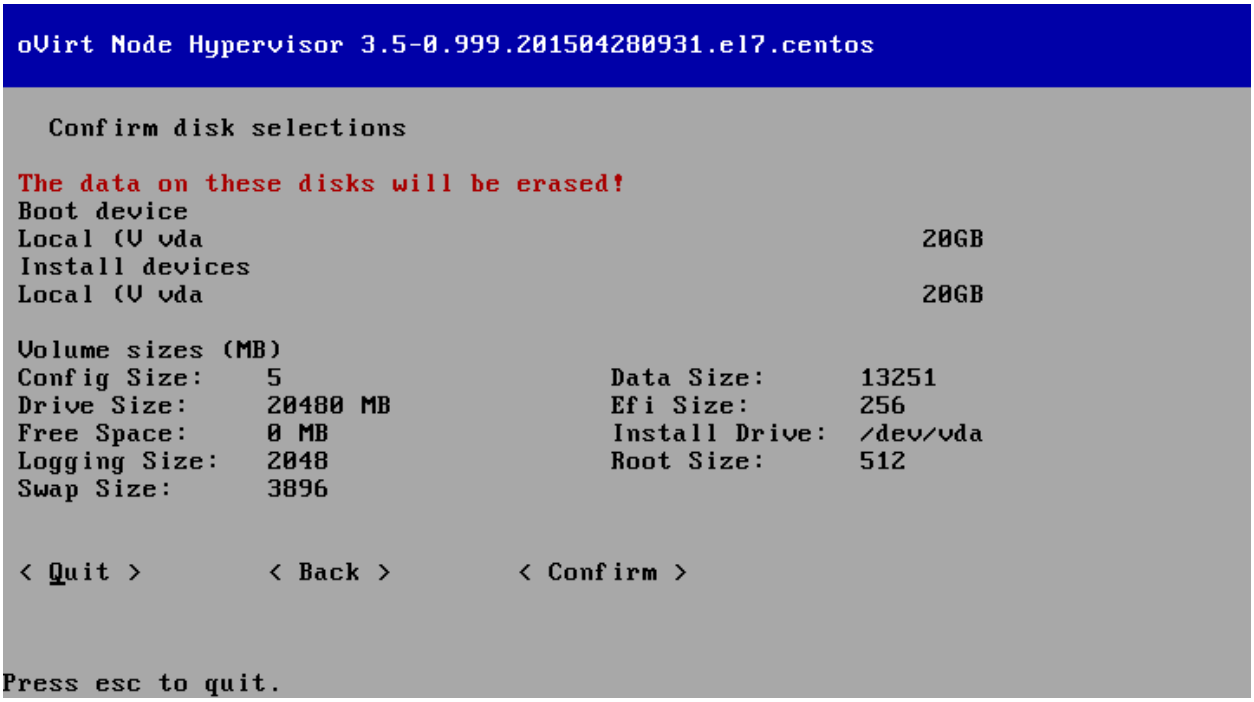
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Confirm disk selections
The data on these disks will be erased!
Boot device
Local (V vda 20GB
Install devices
Local (V vda 20GB
Volume sizes (MB)
Config Size: 5
Drive Size: 20480 MB
Free Space: 0 MB
Logging Size: 2048
Swap Size: 3896
Data Size: 13251
Efi Size: 256
Install Drive: /dev/vda
Root Size: 512
< Quit > < Back > < Confirm >
Press esc to quit.

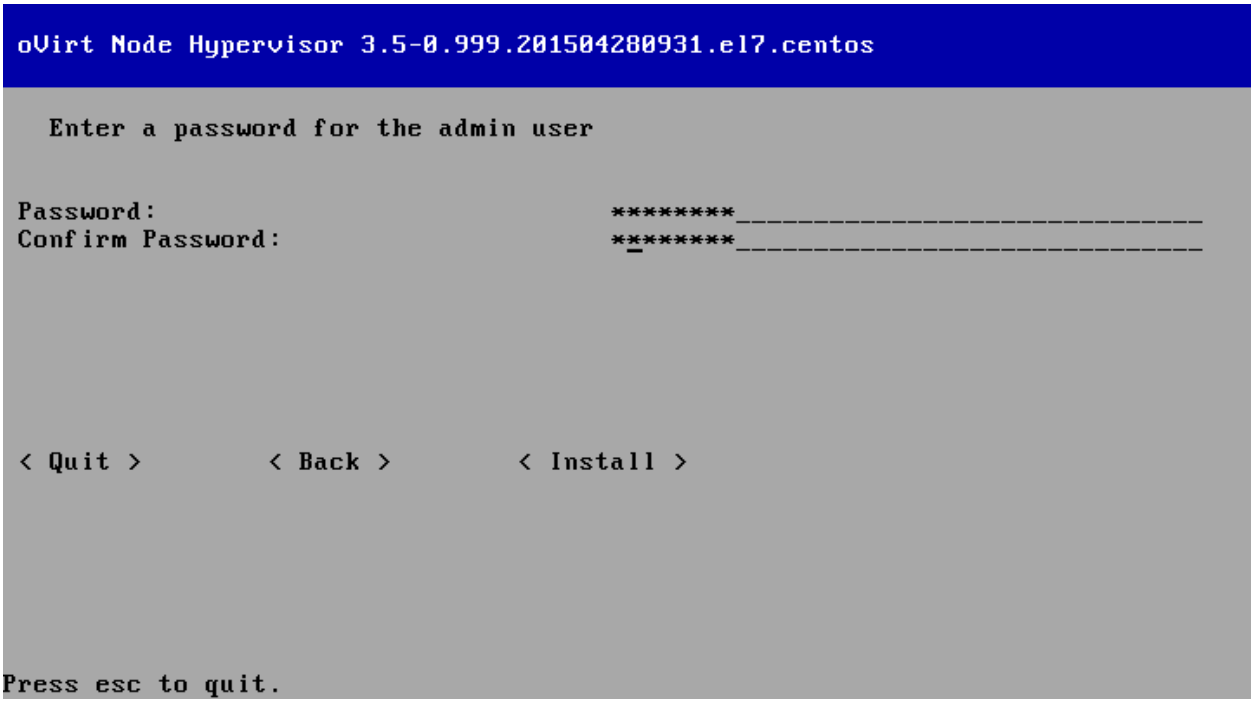
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Enter a password for the admin user
Password:
Confirm Password:
< Quit >
< Back >
< Install >
Press esc to quit.

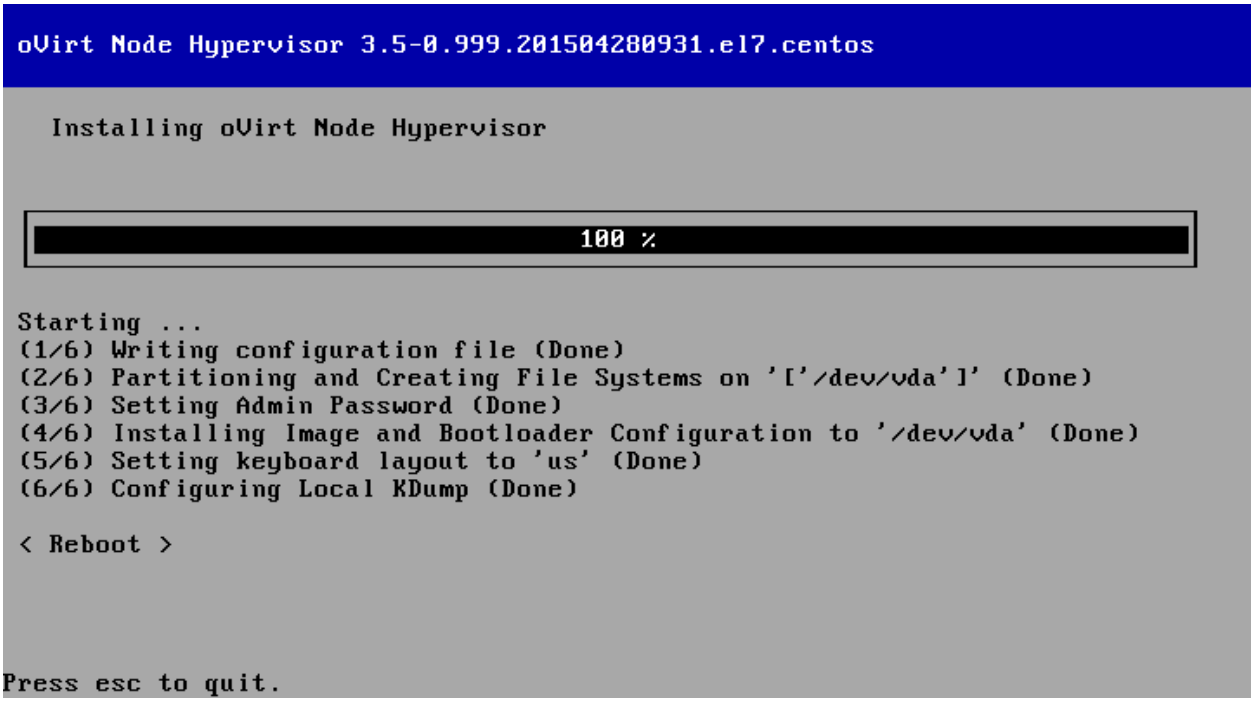
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Installing oVirt Node Hypervisor
100 %
Starting ...
(1/6) Writing configuration file (Done)
(2/6) Partitioning and Creating File Systems on '['/dev/vda']' (Done)
(3/6) Setting Admin Password (Done)
(4/6) Installing Image and Bootloader Configuration to '/dev/vda' (Done)
(5/6) Setting keyboard layout to 'us' (Done)
(6/6) Configuring Local KDump (Done)
< Reboot >
Press esc to quit.

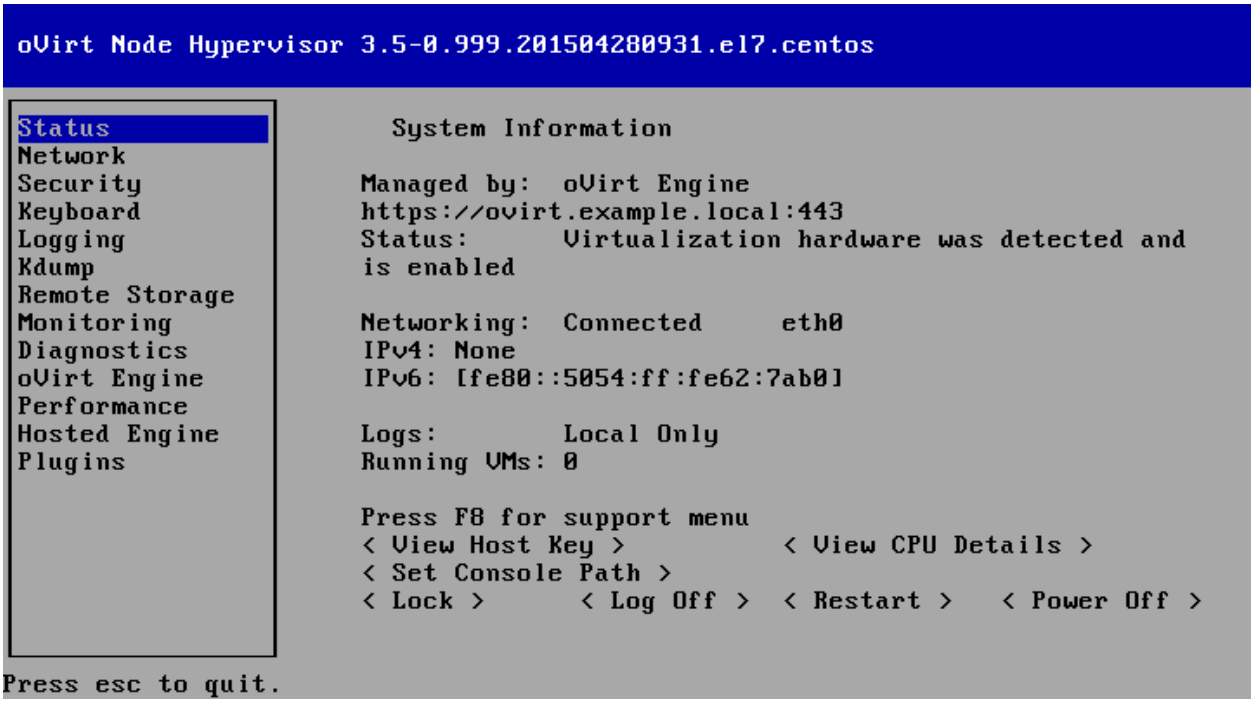
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Status
Network
Security
Keyboard
Logging
Kdump
Remote Storage
Monitoring
Diagnostics
oVirt Engine
Performance
Hosted Engine
Plugins
System Information
Managed by: oVirt Engine
https://ovirt.example.local:443
Status: Virtualization hardware was detected and is enabled
Networking: Connected eth0
IPv4: None
IPv6: [fe80::5054:ff:fe62:7ab0]
Logs: Local Only
Running VMs: 0
Press F8 for support menu
< View Host Key >
< View CPU Details >
< Set Console Path >
< Lock >
< Log Off >
< Restart >
< Power Off >
Press esc to quit.

Chapter 11: Starting Your First Virtual Machine in oVirt

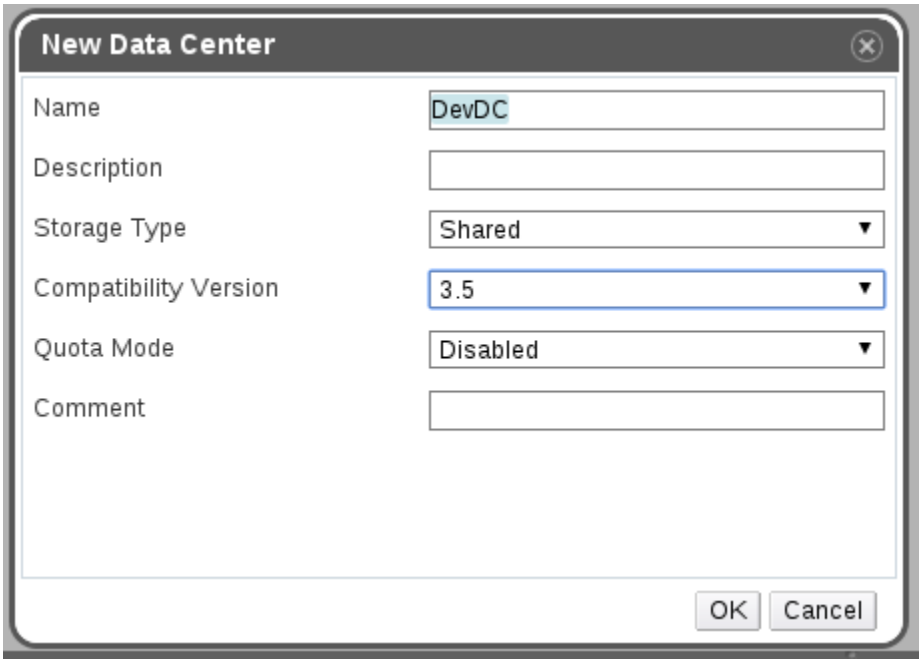

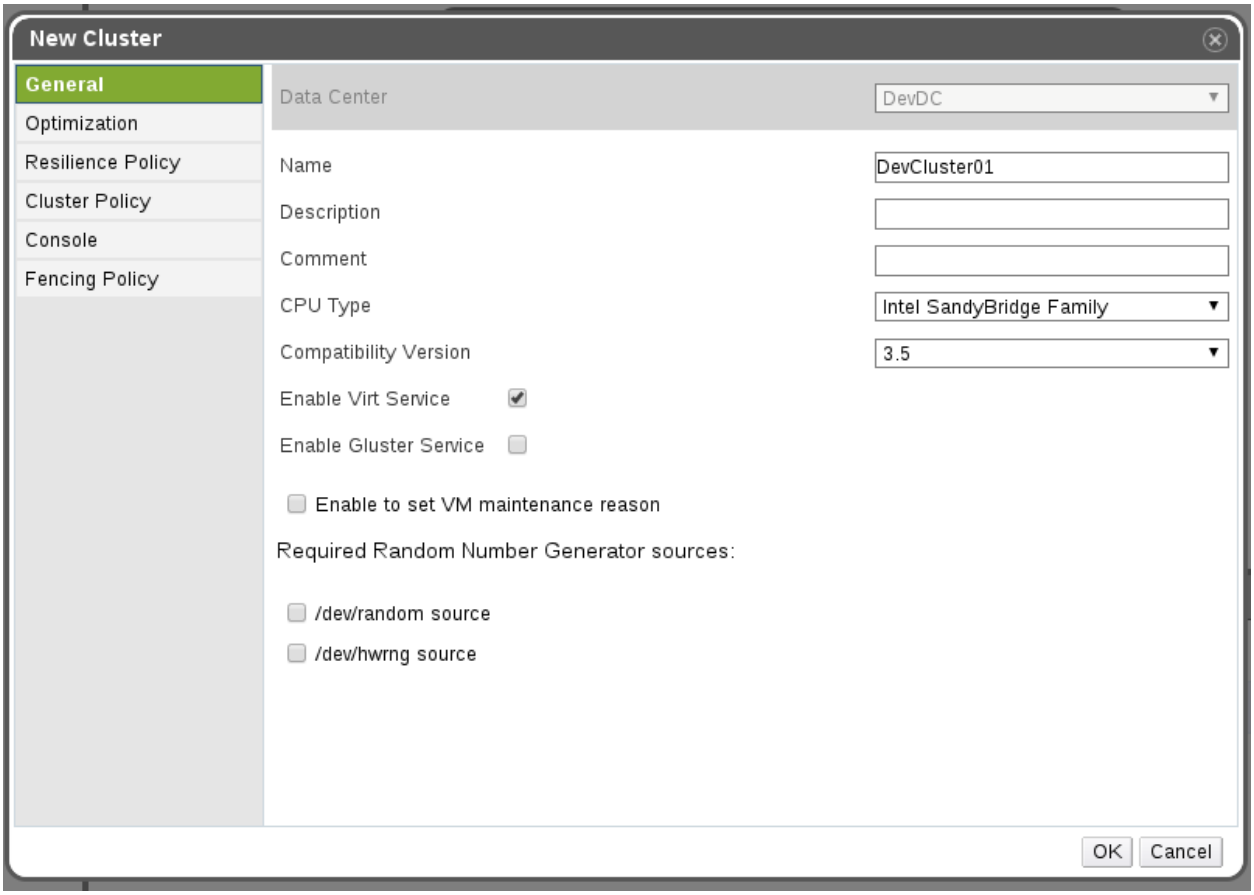

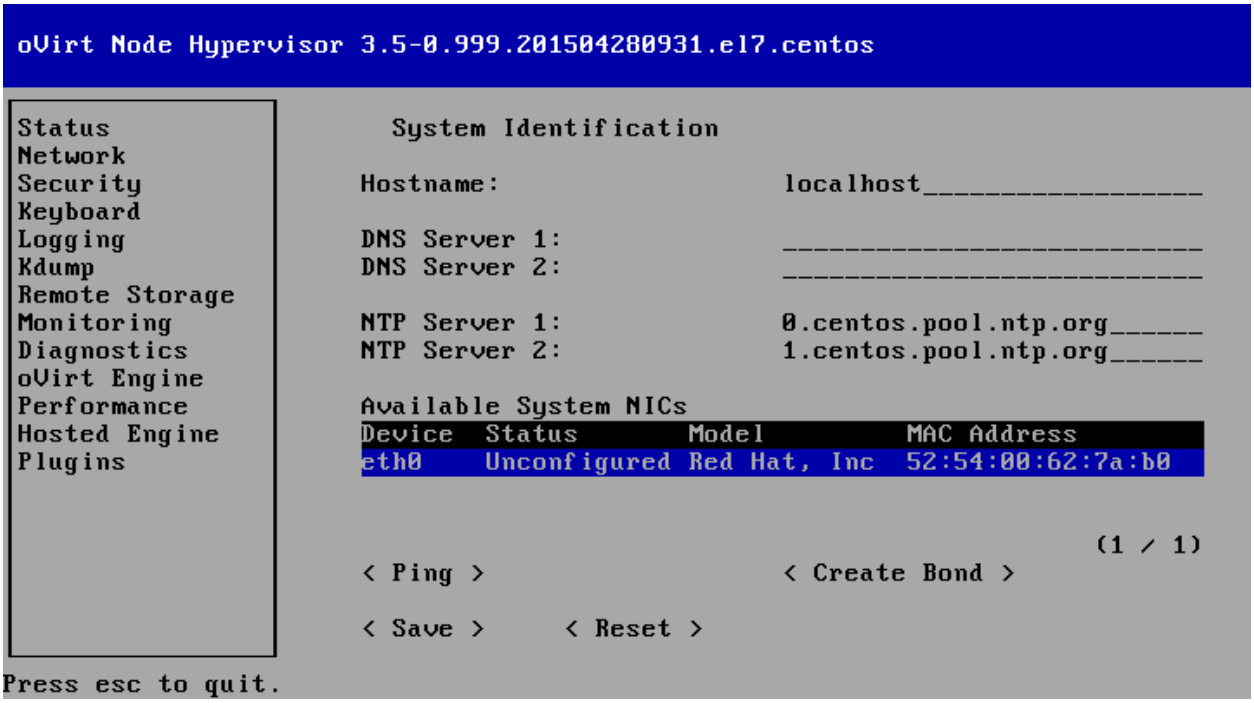
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos
Status
Network
Security
Keyboard
Logging
Kdump
Remote Storage
Monitoring
Diagnostics
oVirt Engine
Performance
Hosted Engine
Plugins
System Identification
Hostname: localhost
DNS Server 1:
DNS Server 2:
NTP Server 1: 0.centos.pool.ntp.org
NTP Server 2: 1.centos.pool.ntp.org
Available System NICs
Device Status Model MAC Address
eth0 Unconfigured Red Hat, Inc 52:54:00:62:7a:b0
(1 / 1)
< Ping >
< Create Bond >
< Save >
< Reset >
Press esc to quit.

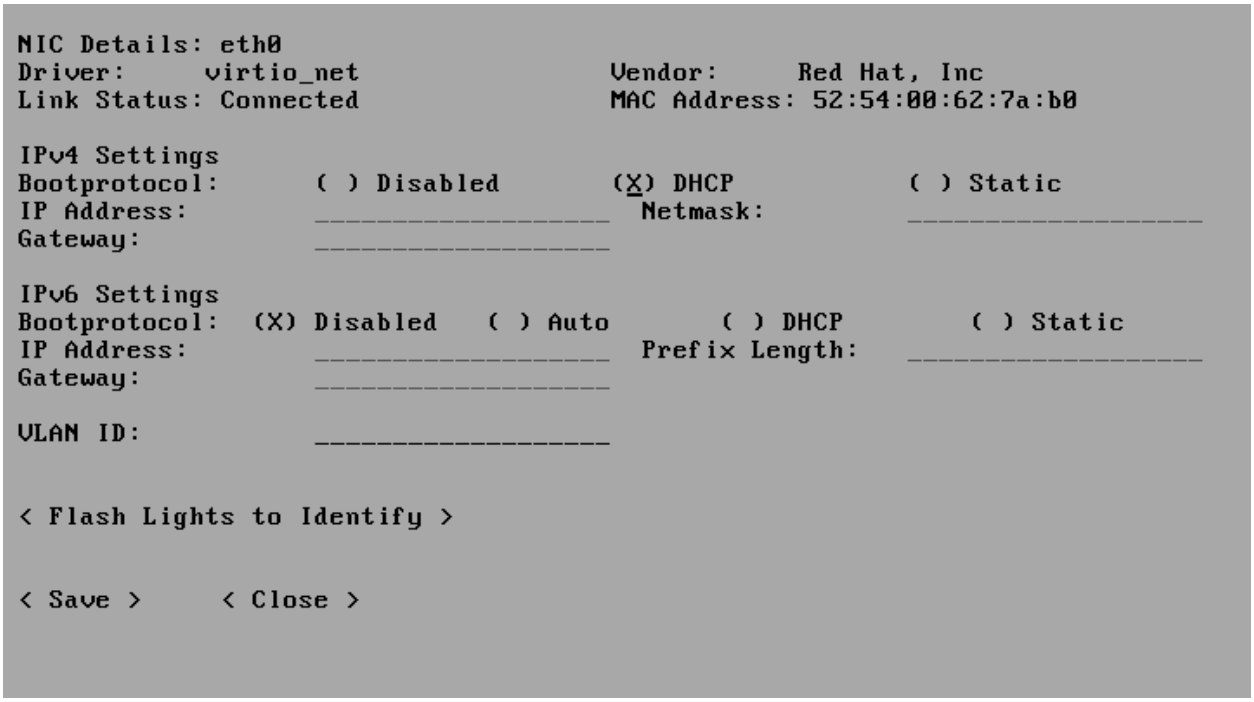
NIC Details: eth0
Driver: virtio_net
Vendor: Red Hat, Inc
Link Status: Connected
MAC Address: 52:54:00:62:7a:b0
IPv4 Settings
Bootprotocol: ( ) Disabled (X) DHCP ( ) Static
IP Address:
Netmask:
Gateway:
IPv6 Settings
Bootprotocol: (X) Disabled ( ) Auto ( ) DHCP ( ) Static
IP Address:
Prefix Length:
Gateway:
VLAN ID:
< Flash Lights to Identify >
< Save >
< Close >

```
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos

Status                    Remote Access
Network
Security                  Enable SSH password          [X]
Keyboard                  authentication
Logging
Kdump
Remote Storage            Strong Random Number Generator
Monitoring
Diagnostics               Disable AES-NI               [ ]
oVirt Engine
Performance               Bytes Used:                  ____________________________
Hosted Engine
Plugins                   Password for the admin user

                          Password:                    ____________________________
                          Confirm Password:            ____________________________

                          < Save >     < Reset >

Press esc to quit.
```

```
oVirt Node Hypervisor 3.5-0.999.201504280931.el7.centos

Status                    oVirt Engine Configuration
Network
Security                  Management Server:           ovirt.example.local________
Keyboard                  Management Server Port:      443________________________
Logging
Kdump                     Certificate Status: N/A
Remote Storage
Monitoring                Optional password for adding Node through oVirt Engine
Diagnostics               Note: Setting password will enable SSH daemon
oVirt Engine              Password:                    ____________________________
Performance               Confirm Password:            ____________________________
Hosted Engine
Plugins



                          < Save & Register >

Press esc to quit.
```

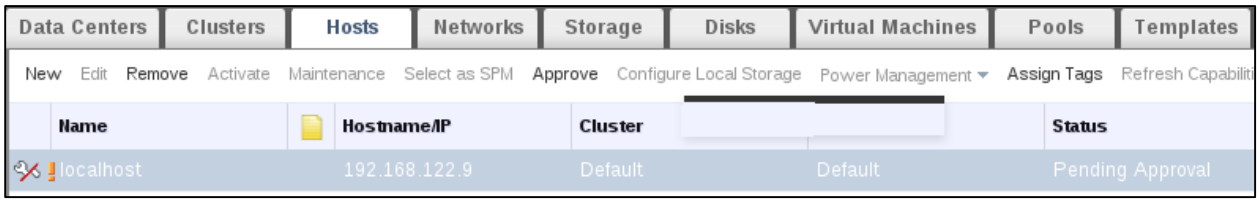

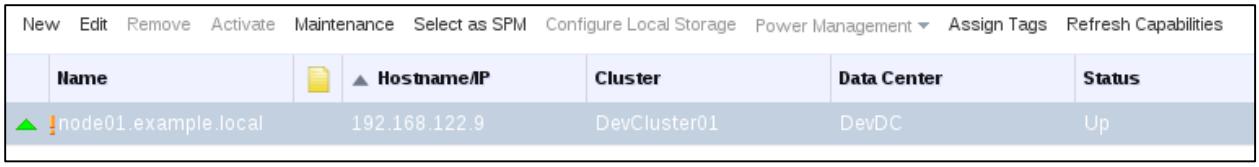

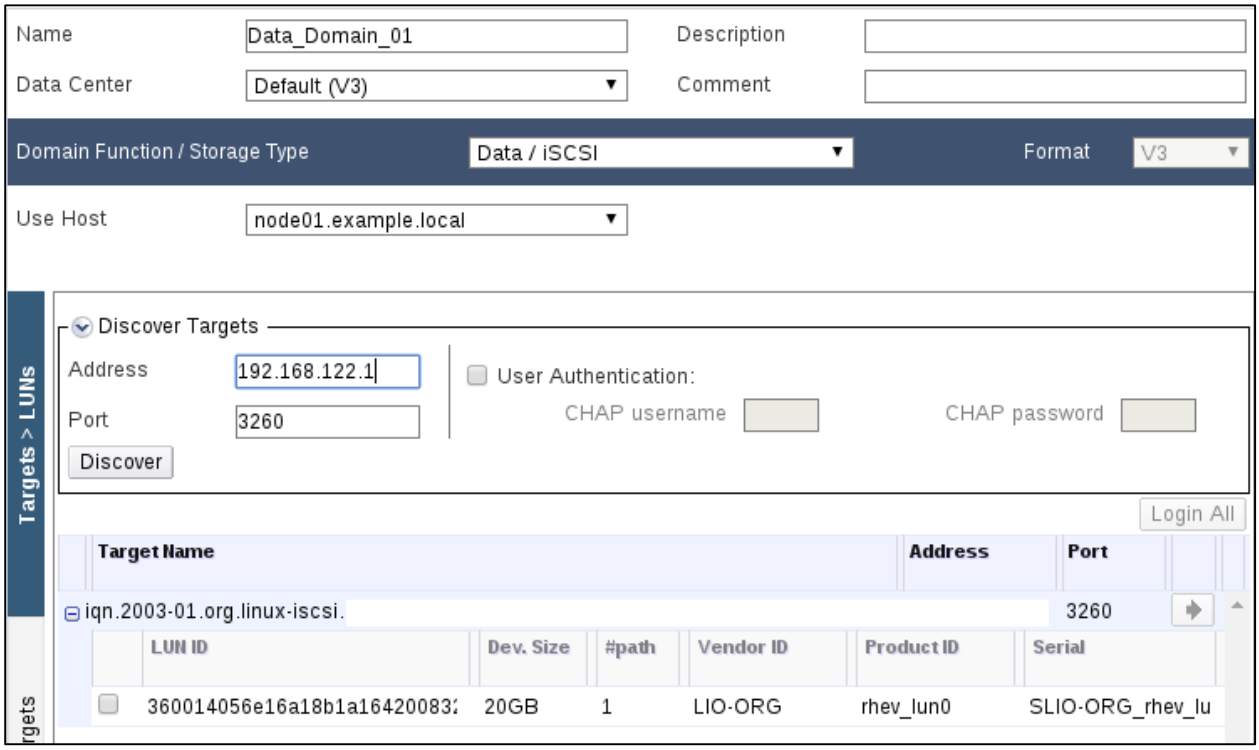

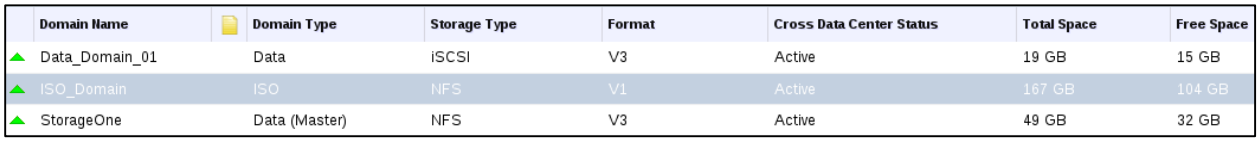


```
360014056e16a18b1a164200832935b1c dm-0 LIO-ORG,rhev_lun0
size=20G features='0' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  `- 2:0:0:0 sda 8:0   active ready running
```

```
  PV                                            VG                                   Fmt  Attr PSize  PFree
  /dev/mapper/360014056e16a18b1a164200832935b1c 4343e974-1359-4b86-bf1f-508dd436504b lvm2 a--  19.62g 15.75g
```

```
4343e974-1359-4b86-bf1f-508dd436504b MDT_CLASS=Data,MDT_DESCRIPTION=Data_Domain_01,
```

```
node01 # vdsClient -s 0 getStorageDomainInfo 4343e974-1359-4b86-bf1f-508dd436504b
        uuid = 4343e974-1359-4b86-bf1f-508dd436504b
        vguuid = dmQ82i-vLkh-q0jZ-jWxQ-tL2w-Au9e-vGawgu
        state = OK
        version = 3
        role = Regular
        type = ISCSI
        class = Data
        pool = ['00000002-0002-0002-0002-0000000001e9']
        name = Data_Domain_01
```

```
node01 # lvs -o name,vg_name
  LV                                   VG
  1efa67e7-d943-4806-8f28-05680dce584a 4343e974-1359-4b86-bf1f-508dd436504b
  58306bf8-418b-40a8-9522-a4d19de7f144 4343e974-1359-4b86-bf1f-508dd436504b
  ids                                  4343e974-1359-4b86-bf1f-508dd436504b
  inbox                                4343e974-1359-4b86-bf1f-508dd436504b
  leases                               4343e974-1359-4b86-bf1f-508dd436504b
  master                               4343e974-1359-4b86-bf1f-508dd436504b
  metadata                             4343e974-1359-4b86-bf1f-508dd436504b
  outbox                               4343e974-1359-4b86-bf1f-508dd436504b
```

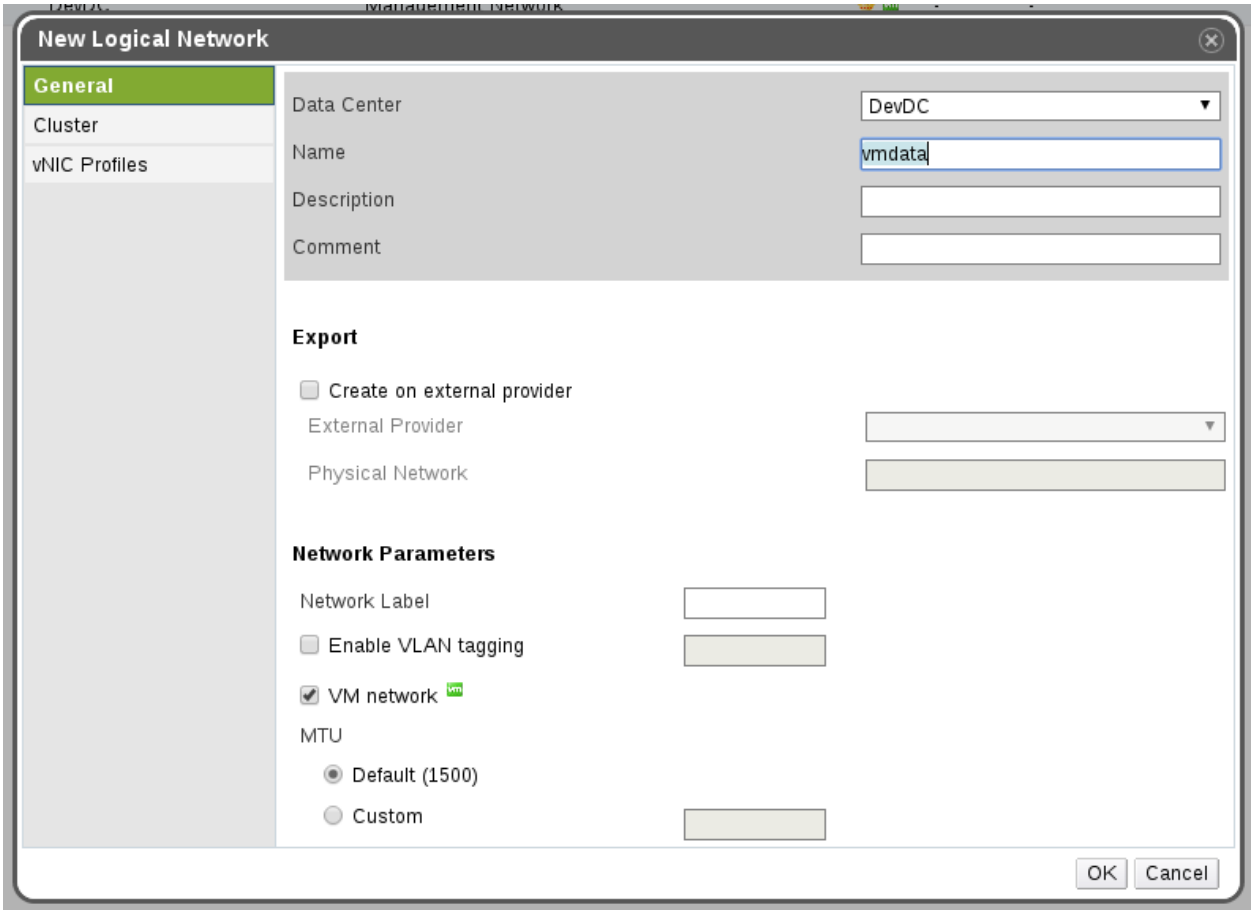


General
Virtual Machines
Network Interfaces
Setup Host Networks
Save Network Configuration

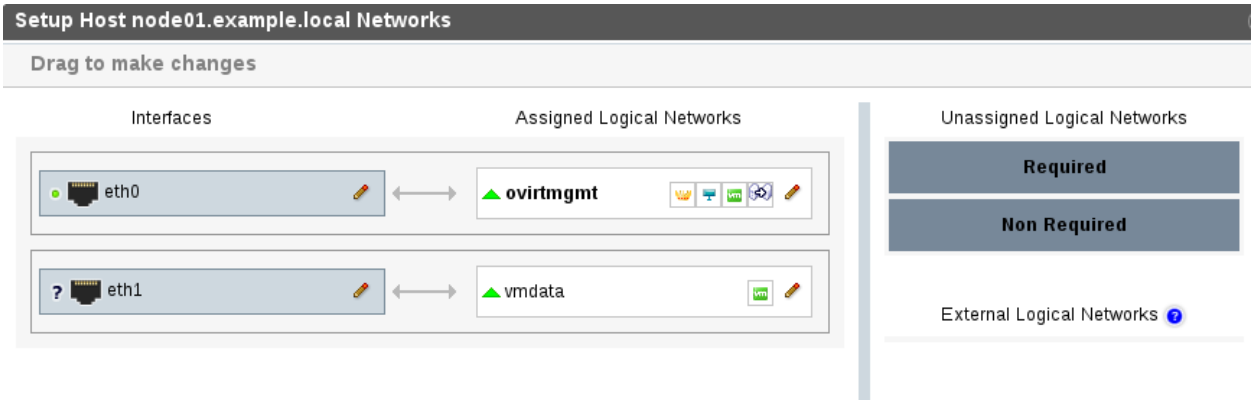
Setup Host node01.example.local Networks
Drag to make changes
Interfaces
Assigned Logical Networks
Unassigned Logical Networks
eth0
ovirtmgmt
Required
eth1
vmdata
Non Required
External Logical Networks

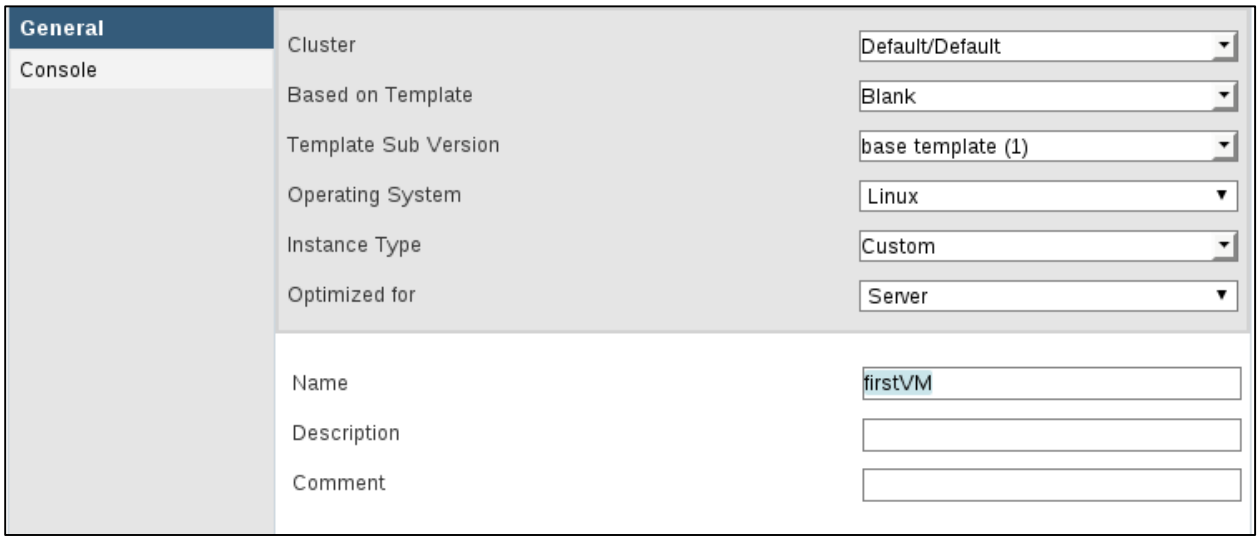
General
Console
Cluster
Default/Default
Based on Template
Blank
Template Sub Version
base template (1)
Operating System
Linux
Instance Type
Custom
Optimized for
Server
Name
firstVM
Description
Comment

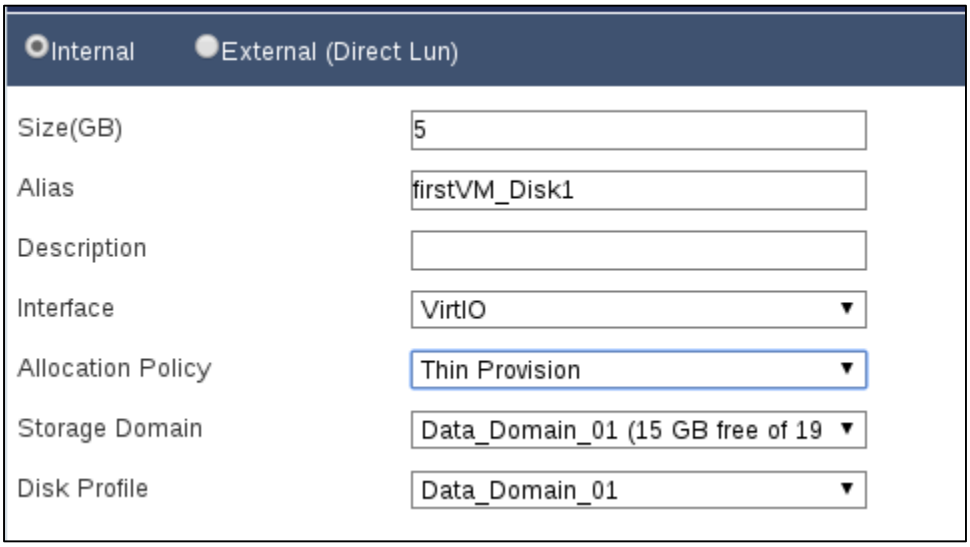
Internal
External (Direct Lun)
Size(GB)
5
Alias
firstVM_Disk1
Description
Interface
VirtIO
Allocation Policy
Thin Provision
Storage Domain
Data_Domain_01 (15 GB free of 19
Disk Profile
Data_Domain_01


```
node01 # virsh -r list
 Id    Name                           State
----------------------------------------------------
 1     firstVM                        running
```

Chapter 12: Deploying OpenStack Private Cloud backed by KVM Virtualization

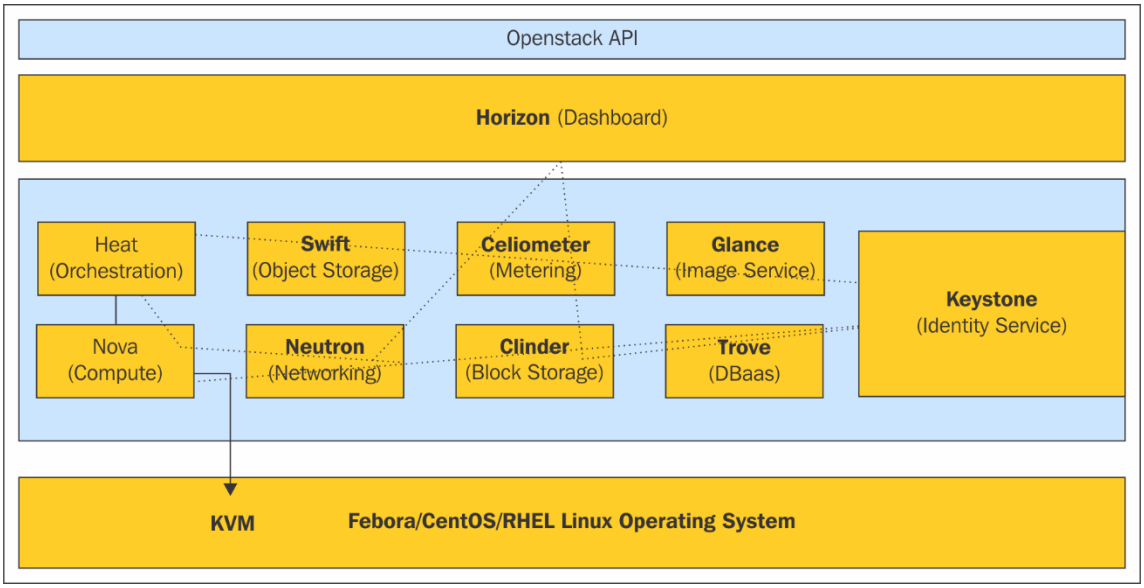


```
[root@dhcp210-192 ~(keystone_admin)]# openstack image list
+--------------------------------------+--------+
| ID                                   | Name   |
+--------------------------------------+--------+
| a106fe87-0201-4d6d-aca3-85bc42b20de7 | cirros |
```

```
[root@openstack ~(keystone_admin)]# glance image-list
+--------------------------------------+----------+
| ID                                   | Name     |
+--------------------------------------+----------+
| a106fe87-0201-4d6d-aca3-85bc42b20de7 | cirros   |
| 78a959eb-7d2e-4eb3-90fd-9234a33fd09a | Fedora22 |
+--------------------------------------+----------+
```

```
[root@openstack ~(keystone_admin)]# nova flavor-list
+----+-----------+-----------+------+-----------+------+-------+-------------+-----------+
| ID | Name      | Memory_MB | Disk | Ephemeral | Swap | VCPUs | RXTX_Factor | Is_Public |
+----+-----------+-----------+------+-----------+------+-------+-------------+-----------+
| 1  | m1.tiny   | 512       | 1    | 0         |      | 1     | 1.0         | True      |
| 2  | m1.small  | 2048      | 20   | 0         |      | 1     | 1.0         | True      |
| 3  | m1.medium | 4096      | 40   | 0         |      | 2     | 1.0         | True      |
| 4  | m1.large  | 8192      | 80   | 0         |      | 4     | 1.0         | True      |
| 5  | m1.xlarge | 16384     | 160  | 0         |      | 8     | 1.0         | True      |
+----+-----------+-----------+------+-----------+------+-------+-------------+-----------+
```

```
[root@openstack ~(keystone_admin)]# nova keypair-list
+------+-------------------------------------------------+
| Name | Fingerprint                                     |
+------+-------------------------------------------------+
| Tom  | 13:2f:c5:fa:ba:29:50:5d:04:18:92:a1:5a:9a:db:e5 |
| John | 12:13:7f:a0:3d:88:d6:fa:d1:7c:4d:00:17:0c:45:b4 |
+------+-------------------------------------------------+
```

```
[root@openstack ~(keystone_admin)]# nova secgroup-list
+--------------------------------------+---------+-------------------------+
| Id                                   | Name    | Description             |
+--------------------------------------+---------+-------------------------+
| 6ca34486-a20b-4da9-8123-185b95d0e5fb | default | Default security group |
+--------------------------------------+---------+-------------------------+
```

```
[root@openstack ~(keystone_admin)]# nova secgroup-create global_web "allow web traffic from the Internet"
+--------------------------------------+------------+---------------------------------+
| Id                                   | Name       | Description                     |
+--------------------------------------+------------+---------------------------------+
| c8c63605-7452-42bf-b2be-3630b81adcba | global_web | allow web traffic from the Internet |
+--------------------------------------+------------+---------------------------------+
```

```
[root@openstack ~(keystone_admin)]# nova secgroup-add-rule global_web tcp 80 80 0.0.0.0/0
+-------------+-----------+---------+-----------+--------------+
| IP Protocol | From Port | To Port | IP Range  | Source Group |
+-------------+-----------+---------+-----------+--------------+
| tcp         | 80        | 80      | 0.0.0.0/0 |              |
+-------------+-----------+---------+-----------+--------------+
```

```
[root@dhcp210-192 ~(keystone_admin)]# nova secgroup-list-rules global_web
+-------------+-----------+---------+-----------+--------------+
| IP Protocol | From Port | To Port | IP Range  | Source Group |
+-------------+-----------+---------+-----------+--------------+
| tcp         | 80        | 80      | 0.0.0.0/0 |              |
| tcp         | 8080      | 8080    | 0.0.0.0/0 |              |
+-------------+-----------+---------+-----------+--------------+
```

```
+--------------------------------------+---------+-----------------------------------------------------+
| id                                   | name    | subnets                                             |
+--------------------------------------+---------+-----------------------------------------------------+
| ef405cc4-dc2a-453e-9df2-e28be91eebe4 | public  | 91be090d-c0d7-423e-8202-0fb29e228b51 172.24.4.224/28 |
| 90905851-38c7-41d4-a331-515c725075ec | private | 63ca5500-0d05-4286-bb7e-d5ab82a499a8 10.0.0.0/24     |
+--------------------------------------+---------+-----------------------------------------------------+
```

```
MariaDB [nova]> SELECT display_name,host,vcpus,memory_mb FROM instances WHERE vm_state='stopped';
+--------------+--------------+-------+-----------+
| display_name | host         | vcpus | memory_mb |
+--------------+--------------+-------+-----------+
| instance033  | compute-node1 |     1 |       512 |
| instance033  | compute-node1 |     1 |       512 |
+--------------+--------------+-------+-----------+
2 rows in set (0.00 sec)
```

Chapter 13: Performance Tuning and Best Practices in KVM

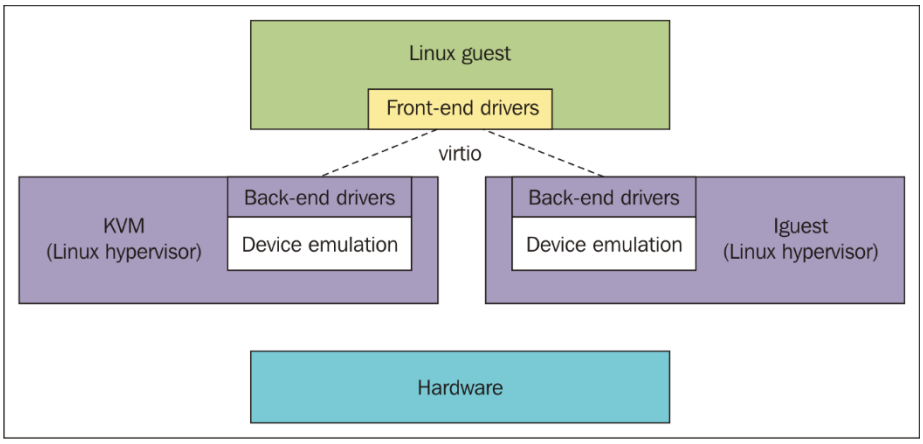

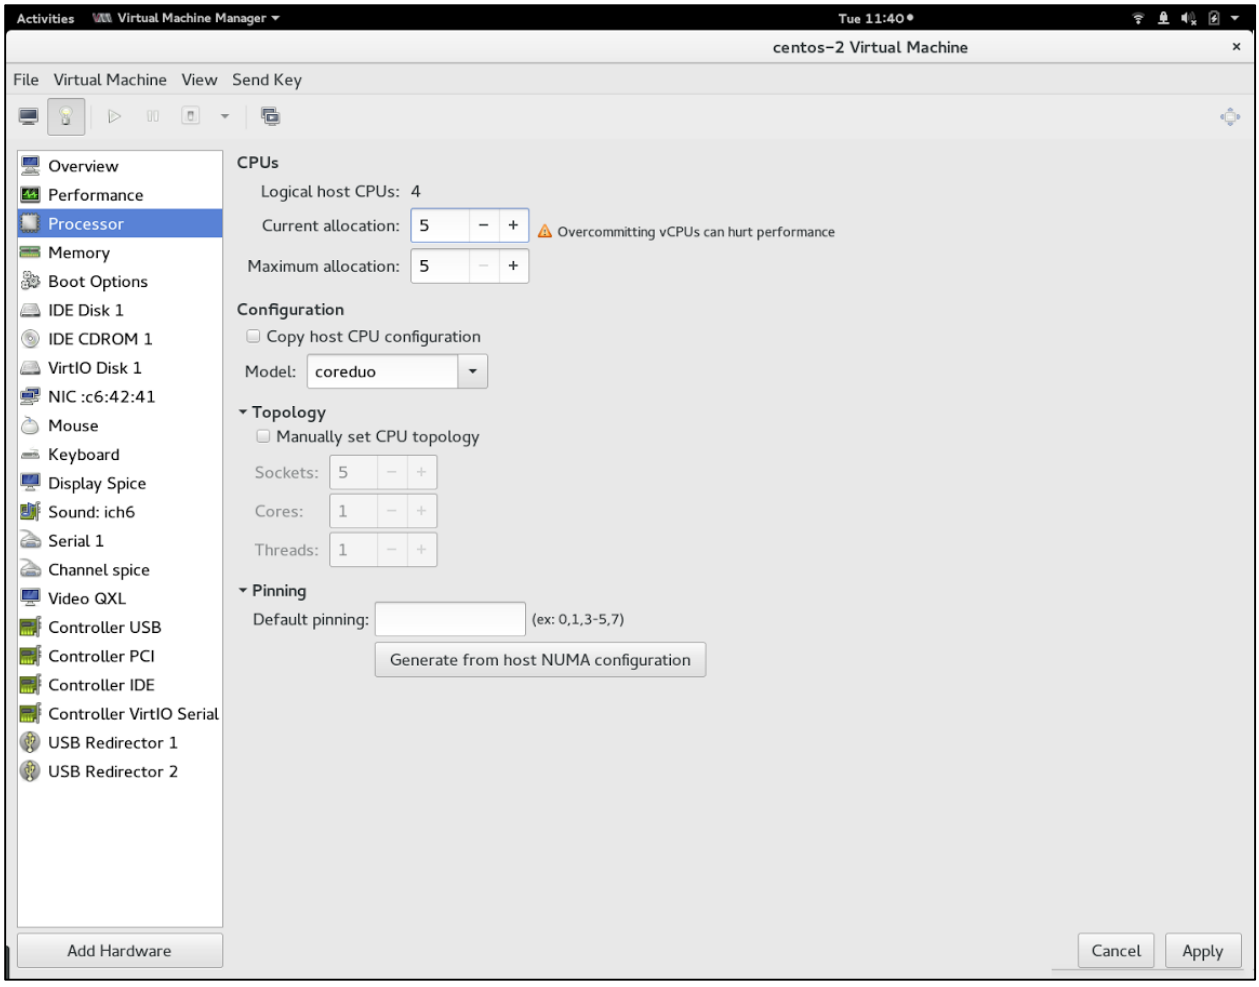

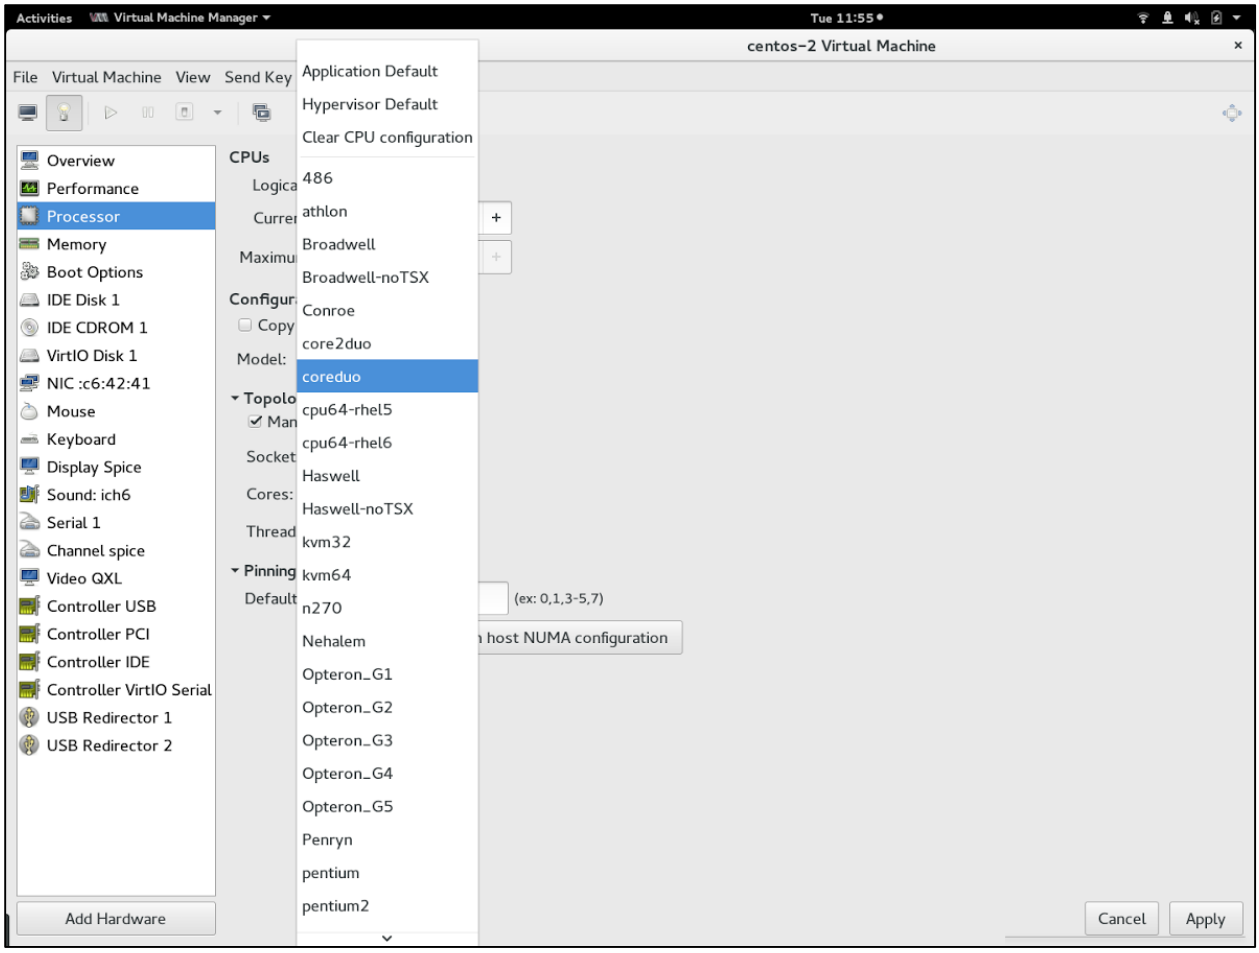
Activities
Virtual Machine Manager
Tue 11:55
centos-2 Virtual Machine
File
Virtual Machine
View
Send Key
Overview
Performance
Processor
Memory
Boot Options
IDE Disk 1
IDE CDROM 1
VirtIO Disk 1
NIC :c6:42:41
Mouse
Keyboard
Display Spice
Sound: ich6
Serial 1
Channel spice
Video QXL
Controller USB
Controller PCI
Controller IDE
Controller VirtIO Serial
USB Redirector 1
USB Redirector 2
Add Hardware
CPUs
Logica
Curre
Maximu
Configur
Copy
Model:
Topolo
Man
Socket
Cores:
Thread
Pinning
Default
Application Default
Hypervisor Default
Clear CPU configuration
486
athlon
Broadwell
Broadwell-noTSX
Conroe
core2duo
coreduo
cpu64-rhel5
cpu64-rhel6
Haswell
Haswell-noTSX
kvm32
kvm64
n270
Nehalem
Opteron_G1
Opteron_G2
Opteron_G3
Opteron_G4
Opteron_G5
Penryn
pentium
pentium2
(ex: 0,1,3-5,7)
host NUMA configuration
Cancel
Apply

```
humble-lap $ virsh capabilities
<capabilities>

  <host>
    <uuid>bf4a43ce-e99c-4bc5-b2f2-8aacfcfcd552</uuid>
    <cpu>
      <arch>x86_64</arch>
      <model>Broadwell</model>
      <vendor>Intel</vendor>
      <topology sockets='1' cores='2' threads='2'/>
      <feature name='invtsc'/>
      <feature name='abm'/>
      <feature name='pdpe1gb'/>
      <feature name='rdrand'/>
      <feature name='f16c'/>
      <feature name='osxsave'/>
      <feature name='pdcm'/>
      <feature name='xtpr'/>
      <feature name='tm2'/>
      <feature name='est'/>
      <feature name='smx'/>
      <feature name='vmx'/>
      <feature name='ds_cpl'/>
      <feature name='monitor'/>
      <feature name='dtes64'/>
      <feature name='pbe'/>
      <feature name='tm'/>
      <feature name='ht'/>
      <feature name='ss'/>
      <feature name='acpi'/>
      <feature name='ds'/>
      <feature name='vme'/>
      <pages unit='KiB' size='4'/>
      <pages unit='KiB' size='2048'/>
    </cpu>
```

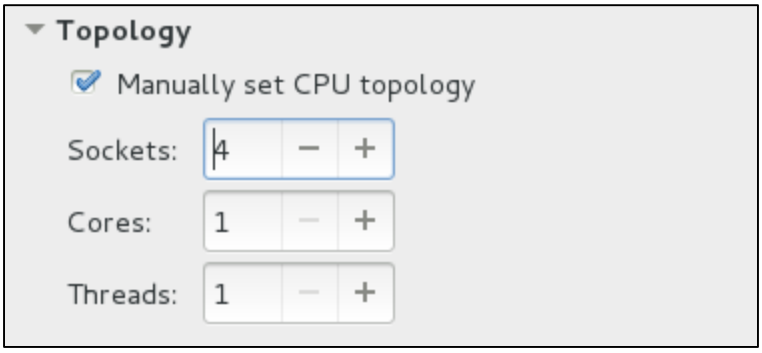

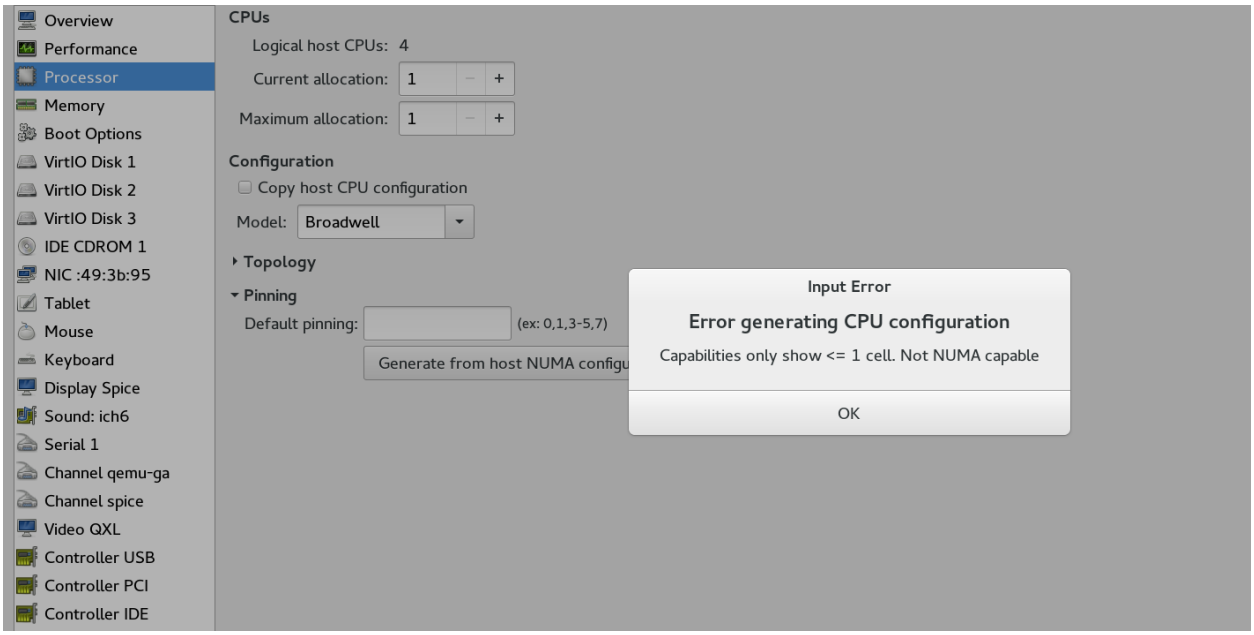


```
 humble-lap $ numactl --hardware
available: 1 nodes (0)
node 0 cpus: 0 1 2 3
node 0 size: 7668 MB
node 0 free: 673 MB
node distances:
node   0
  0:  10
```

```
[humble-lap ]$ virsh vcpupin --help
  NAME
    vcpupin - control or query domain vcpu affinity

  SYNOPSIS
    vcpupin <domain> [--vcpu <number>] [--cpulist <string>] [--config] [--live] [--current]

  DESCRIPTION
    Pin domain VCPUs to host physical CPUs.

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    --vcpu <number>  vcpu number
    --cpulist <string>  host cpu number(s) to set, or omit option to query
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain


[humble-lap ]$
```

```
 humble-lap $ virsh nodeinfo
CPU model:           x86_64
CPU(s):              4
CPU frequency:       2895 MHz
CPU socket(s):       1
Core(s) per socket:  2
Thread(s) per core:  2
NUMA cell(s):        1
Memory size:         7852832 KiB
```

```
<topology>
  <cells num='1'>
    <cell id='0'>
      <memory unit='KiB'>7852832</memory>
      <pages unit='KiB' size='4'>1963208</pages>
      <pages unit='KiB' size='2048'>0</pages>
      <distances>
        <sibling id='0' value='10'/>
      </distances>
      <cpus num='4'>
        <cpu id='0' socket_id='0' core_id='0' siblings='0-1'/>
        <cpu id='1' socket_id='0' core_id='0' siblings='0-1'/>
        <cpu id='2' socket_id='0' core_id='1' siblings='2-3'/>
        <cpu id='3' socket_id='0' core_id='1' siblings='2-3'/>
      </cpus>
    </cell>
  </cells>
</topology>
```

```
humble-lap $  virsh vcpupin centos-1
VCPU: CPU Affinity
----------------------------------
   0: 0-3
   1: 0-3
   2: 0-3
   3: 0-3

humble-lap $
```

```
 humble-lap $ virsh vcpupin centos-1 1 3

 humble-lap $ virsh vcpupin centos-1
VCPU: CPU Affinity
----------------------------------
   0: 0-3
   1: 3
   2: 0-3
   3: 0-3

 humble-lap $ virsh dumpxml centos-1
<domain type='kvm' id='2'>
  <name>centos-1</name>
  <uuid>4fe91628-1bdc-4765-8f01-d6387400c4c9</uuid>
  <memory unit='KiB'>2097152</memory>
  <currentMemory unit='KiB'>2097152</currentMemory>
  <vcpu placement='static'>4</vcpu>
  <cputune>
    <vcpupin vcpu='1' cpuset='3'/>
  </cputune>
```

```
humble-lap $ virsh vcpupin centos-1 0 3

humble-lap $ virsh vcpupin centos-1 1 2

humble-lap $ virsh vcpupin centos-1 2 3

humble-lap $ virsh vcpupin centos-1 3 1

humble-lap $ virsh vcpupin centos-1
VCPU: CPU Affinity
----------------------------------
  0: 3
  1: 2
  2: 3
  3: 1
```

```
 humble-lap $ virsh vcpuinfo
error: command 'vcpuinfo' requires <domain> option
 humble-lap $ virsh vcpuinfo centos-1
VCPU:           0
CPU:            3
State:          running
CPU time:       8.7s
CPU Affinity:   ---y

VCPU:           1
CPU:            2
State:          running
CPU time:       7.2s
CPU Affinity:   --y-

VCPU:           2
CPU:            3
State:          running
CPU time:       9.9s
CPU Affinity:   ---y

VCPU:           3
CPU:            1
State:          running
CPU time:       5.6s
CPU Affinity:   -y--
```

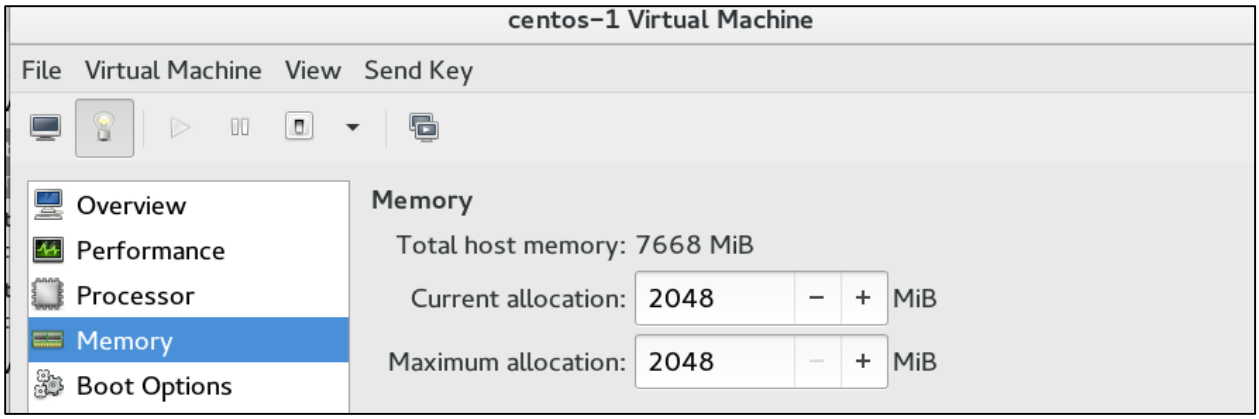

```
[ humble-lap ]$ virsh setmem --help
  NAME
    setmem - change memory allocation

  SYNOPSIS
    setmem <domain> <size> [--config] [--live] [--current]

  DESCRIPTION
    Change the current memory allocation in the guest domain.

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    [--size] <number>  new memory size, as scaled integer (default KiB)
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain


[ humble-lap ]$ virsh setmaxmem --help
  NAME
    setmaxmem - change maximum memory limit

  SYNOPSIS
    setmaxmem <domain> <size> [--config] [--live] [--current]

  DESCRIPTION
    Change the maximum memory allocation limit in the guest domain.

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    [--size] <number>  new maximum memory size, as scaled integer (default KiB)
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain
```

```
[humble-lap ]$ virsh list
 Id    Name                           State
----------------------------------------------------
 2     centos-1                       running

[humble-lap ]$ virsh memtune centos-1
hard_limit     : 9007199254740988
soft_limit     : 9007199254740988
swap_hard_limit: 9007199254740988

[humble-lap ]$
```

```
humble-lap $  virsh help memtune
  NAME
    memtune - Get or set memory parameters

  SYNOPSIS
    memtune <domain> [--hard-limit <number>] [--soft-limit <number>]
 [--swap-hard-limit <number>] [--min-guarantee <number>] [--config]
[--live] [--current]

  DESCRIPTION
    Get or set the current memory parameters for a guest domain.
    To get the memory parameters use following command:

    virsh # memtune <domain>

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    --hard-limit <number>  Max memory, as scaled integer (default Ki
B)
    --soft-limit <number>  Memory during contention, as scaled integ
er (default KiB)
    --swap-hard-limit <number>  Max memory plus swap, as scaled inte
ger (default KiB)
    --min-guarantee <number>  Min guaranteed memory, as scaled integ
er (default KiB)
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain


humble-lap $
```

```
[ humble-lap ]$ cat /proc/meminfo |grep -i huge
AnonHugePages:         0 kB
HugePages_Total:       0
HugePages_Free:        0
HugePages_Rsvd:        0
HugePages_Surp:        0
Hugepagesize:       2048 kB
[ humble-lap ]$
```

```
vm.hugepages_treat_as_movable = 0
vm.hugetlb_shm_group = 0
vm.nr_hugepages = 0
vm.nr_hugepages_mempolicy = 0
vm.nr_overcommit_hugepages = 0
```

```
[ humble-lap ]$ mount |grep huge
cgroup on /sys/fs/cgroup/hugetlb type cgroup (rw,nosuid,nodev,noexec,relatime,hugetlb)
hugetlbfs on /dev/hugepages type hugetlbfs (rw,relatime,seclabel)
[ humble-lap ]$ cd /dev/hugepages/
[ humble-lap ]$ ls
libvirt
[ humble-lap ]$ cd libvirt/
[ humble-lap ]$ ls
qemu
```

```
 humble-lap $ cat /etc/redhat-release
Fedora release 22 (Twenty Two)
 humble-lap $ uname -r
4.2.3-200.fc22.x86_64
 humble-lap $ cat /boot/config-4.2.3-200.fc22.x86_64 |grep CONFIG_KSM
CONFIG_KSH=y
 humble-lap $ dnf install ksm
```

```
 humble-lap $ rpm -ql ksm
/etc/ksmtuned.conf
/etc/sysconfig/ksm
/usr/lib/systemd/system/ksm.service
/usr/lib/systemd/system/ksmtuned.service
/usr/libexec/ksmctl
/usr/sbin/ksmtuned
 humble-lap $
```

```
 humble-lap $  ls /sys/kernel/mm/ksm/*
/sys/kernel/mm/ksm/full_scans
/sys/kernel/mm/ksm/merge_across_nodes
/sys/kernel/mm/ksm/pages_shared
/sys/kernel/mm/ksm/pages_sharing
/sys/kernel/mm/ksm/pages_to_scan
/sys/kernel/mm/ksm/pages_unshared
/sys/kernel/mm/ksm/pages_volatile
/sys/kernel/mm/ksm/run
/sys/kernel/mm/ksm/sleep_millisecs
 humble-lap $
```

```
[humble-lap ]$ systemctl status ksm
● ksm.service - Kernel Samepage Merging
   Loaded: loaded (/usr/lib/systemd/system/ksm.service; enabled; vendor preset: enabled)
   Active: active (exited) since Fri 2015-11-20 14:57:49 IST; 3h 21min ago
  Process: 1327 ExecStart=/usr/libexec/ksmctl start (code=exited, status=0/SUCCESS)
 Main PID: 1327 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/ksm.service

Warning: Journal has been rotated since unit was started. Log output is incomplete or unavailable.
[humble-lap ]$ ps aux |grep ksm
root        39  0.0  0.0      0     0 ?        SN   14:56   0:00 [ksmd]
```

```
 humble-lap $ cat /sys/kernel/mm/ksm/*
9
1
1084
24585
100
118769
1338
1
20
 humble-lap $ cat /sys/kernel/mm/ksm/*
9
1
1084
24585
100
118742
1365
1
20
 humble-lap $ cat /sys/kernel/mm/ksm/*
9
1
1084
24585
100
118741
1366
1
20
 humble-lap $ cat /sys/kernel/mm/ksm/full_scans
10
 humble-lap $ cat /sys/kernel/mm/ksm/*
10
1
1084
24584
100
118714
1394
1
20
 humble-lap $
```

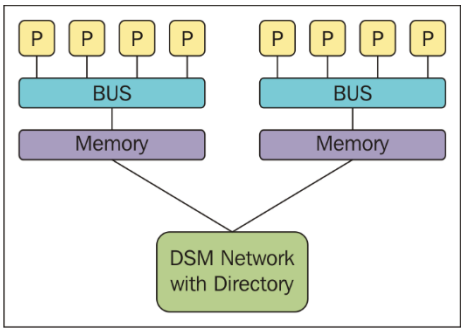


```
[ humble-numaserver ]$ numactl -H
available: 2 nodes (0-1)
node 0 cpus: 0 2 4 6 8 10 12 14 16 18 20 22 24 26 28 30
node 0 size: 131026 MB
node 0 free: 114933 MB
node 1 cpus: 1 3 5 7 9 11 13 15 17 19 21 23 25 27 29 31
node 1 size: 131072 MB
node 1 free: 112458 MB
node distances:
node   0   1
  0:  10  20
  1:  20  10
[ humble-numaserver ]$
```

```
<topology>
  <cells num='2'>
    <cell id='0'>
      <memory unit='KiB'>134171180</memory>
      <pages unit='KiB' size='4'>33542795</pages>
      <pages unit='KiB' size='2048'>0</pages>
      <distances>
        <sibling id='0' value='10'/>
        <sibling id='1' value='20'/>
      </distances>
      <cpus num='16'>
        <cpu id='0' socket_id='0' core_id='0' siblings='0,16'/>
        <cpu id='2' socket_id='0' core_id='1' siblings='2,18'/>
        <cpu id='4' socket_id='0' core_id='2' siblings='4,20'/>
        <cpu id='6' socket_id='0' core_id='3' siblings='6,22'/>
        <cpu id='8' socket_id='0' core_id='4' siblings='8,24'/>
        <cpu id='10' socket_id='0' core_id='5' siblings='10,26'/>
        <cpu id='12' socket_id='0' core_id='6' siblings='12,28'/>
        <cpu id='14' socket_id='0' core_id='7' siblings='14,30'/>
        <cpu id='16' socket_id='0' core_id='0' siblings='0,16'/>
        <cpu id='18' socket_id='0' core_id='1' siblings='2,18'/>
        <cpu id='20' socket_id='0' core_id='2' siblings='4,20'/>
        <cpu id='22' socket_id='0' core_id='3' siblings='6,22'/>
        <cpu id='24' socket_id='0' core_id='4' siblings='8,24'/>
        <cpu id='26' socket_id='0' core_id='5' siblings='10,26'/>
        <cpu id='28' socket_id='0' core_id='6' siblings='12,28'/>
        <cpu id='30' socket_id='0' core_id='7' siblings='14,30'/>
      </cpus>
    </cell>
    <cell id='1'>
      <memory unit='KiB'>134217728</memory>
      <pages unit='KiB' size='4'>33554432</pages>
      <pages unit='KiB' size='2048'>0</pages>
      <distances>
        <sibling id='0' value='20'/>
        <sibling id='1' value='10'/>
      </distances>
      <cpus num='16'>
        <cpu id='1' socket_id='1' core_id='0' siblings='1,17'/>
        <cpu id='3' socket_id='1' core_id='1' siblings='3,19'/>
        <cpu id='5' socket_id='1' core_id='2' siblings='5,21'/>
        <cpu id='7' socket_id='1' core_id='3' siblings='7,23'/>
        <cpu id='9' socket_id='1' core_id='4' siblings='9,25'/>
        <cpu id='11' socket_id='1' core_id='5' siblings='11,27'/>
        <cpu id='13' socket_id='1' core_id='6' siblings='13,29'/>
        <cpu id='15' socket_id='1' core_id='7' siblings='15,31'/>
        <cpu id='17' socket_id='1' core_id='0' siblings='1,17'/>
        <cpu id='19' socket_id='1' core_id='1' siblings='3,19'/>
        <cpu id='21' socket_id='1' core_id='2' siblings='5,21'/>
        <cpu id='23' socket_id='1' core_id='3' siblings='7,23'/>
        <cpu id='25' socket_id='1' core_id='4' siblings='9,25'/>
        <cpu id='27' socket_id='1' core_id='5' siblings='11,27'/>
        <cpu id='29' socket_id='1' core_id='6' siblings='13,29'/>
        <cpu id='31' socket_id='1' core_id='7' siblings='15,31'/>
      </cpus>
    </cell>
  </cells>
</topology>
```

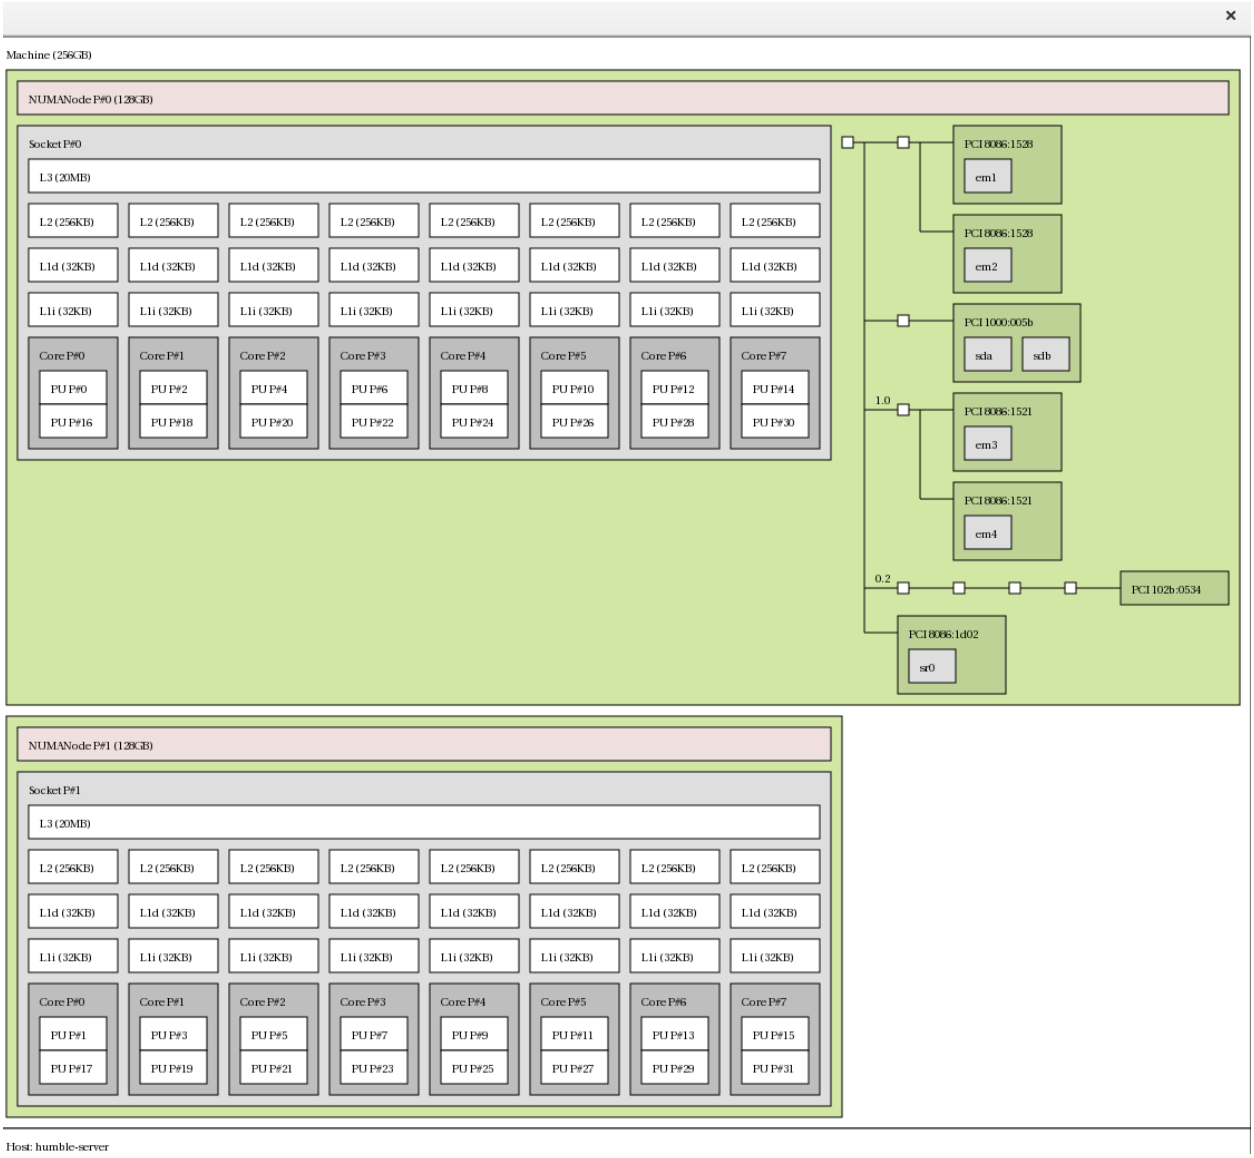
Machine (256GB)
NUMANode P#0 (128GB)
Socket P#0
L3 (20MB)
L2 (256KB)
L1d (32KB)
L1i (32KB)
Core P#0
PU P#0
PU P#16
Core P#1
PU P#2
PU P#18
Core P#2
PU P#4
PU P#20
Core P#3
PU P#6
PU P#22
Core P#4
PU P#8
PU P#24
Core P#5
PU P#10
PU P#26
Core P#6
PU P#12
PU P#28
Core P#7
PU P#14
PU P#30
PCI 8086:1528
em1
PCI 8086:1528
em2
PCI 1000:005b
sda
sdb
1.0
PCI 8086:1521
em3
PCI 8086:1521
em4
0.2
PCI 102b:0534
PCI 8086:1d02
sr0
NUMANode P#1 (128GB)
Socket P#1
L3 (20MB)
L2 (256KB)
L1d (32KB)
L1i (32KB)
Core P#0
PU P#1
PU P#17
Core P#1
PU P#3
PU P#19
Core P#2
PU P#5
PU P#21
Core P#3
PU P#7
PU P#23
Core P#4
PU P#9
PU P#25
Core P#5
PU P#11
PU P#27
Core P#6
PU P#13
PU P#29
Core P#7
PU P#15
PU P#31
Host: humble-server

```
[humble-lap ]$ virsh numatune --help
  NAME
    numatune - Get or set numa parameters

  SYNOPSIS
    numatune <domain> [--mode <string>] [--nodeset <string>] [--config] [--live] [--current]

  DESCRIPTION
    Get or set the current numa parameters for a guest domain.
    To get the numa parameters use following command:

    virsh # numatune <domain>

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    --mode <string>  NUMA mode, one of strict, preferred and interleave
or a number from the virDomainNumatuneMemMode enum
    --nodeset <string>  NUMA node selections to set
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain


[humble-lap ]$
```

```
[ humble-numaserver ]$ numastat -c qemu-kvm

Per-node process memory usage (in MBs)
PID              Node 0 Node 1 Total
---------------  ------ ------ -----
1479 (qemu-kvm)      10   8865  8875
2119 (qemu-kvm)     667     77   744
2194 (qemu-kvm)    1465      0  1465
18404 (qemu-kvm)     30     25    54
20129 (qemu-kvm)   2182      0  2182
32548 (qemu-kvm)     34     16    50
---------------  ------ ------ -----
Total              4389   8982 13371
```

Readonly:

Shareable:

▾ Advanced options

Disk bus: VirtIO

Serial number:

Storage format: qcow2

▾ Performance options

Cache mode: Hypervisor default

IO mode: Hypervisor default

▾ IO Tuning

	KiBytes/Sec	IOPS/Sec
Read:	0	0
Write:	0	0
Total:	0	0

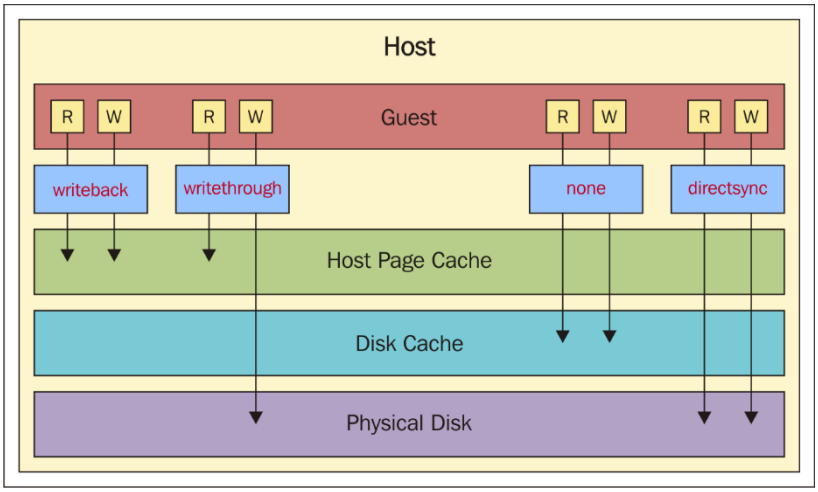

```
[humble-lap ]$ virsh blkdeviotune --help
  NAME
    blkdeviotune - Set or query a block device I/O tuning parameters.

  SYNOPSIS
    blkdeviotune <domain> <device> [--total-bytes-sec <number>] [--read
-bytes-sec <number>] [--write-bytes-sec <number>] [--total-iops-sec <nu
mber>] [--read-iops-sec <number>] [--write-iops-sec <number>] [--total-
bytes-sec-max <number>] [--read-bytes-sec-max <number>] [--write-bytes-
sec-max <number>] [--total-iops-sec-max <number>] [--read-iops-sec-max
<number>] [--write-iops-sec-max <number>] [--size-iops-sec <number>] [-
-config] [--live] [--current]

  DESCRIPTION
    Set or query disk I/O parameters such as block throttling.

  OPTIONS
    [--domain] <string>  domain name, id or uuid
    [--device] <string>  block device
    --total-bytes-sec <number>  total throughput limit in bytes per sec
ond
    --read-bytes-sec <number>  read throughput limit in bytes per secon
d
    --write-bytes-sec <number>  write throughput limit in bytes per sec
ond
    --total-iops-sec <number>  total I/O operations limit per second
    --read-iops-sec <number>  read I/O operations limit per second
    --write-iops-sec <number>  write I/O operations limit per second
    --total-bytes-sec-max <number>  total max in bytes
    --read-bytes-sec-max <number>  read max in bytes
    --write-bytes-sec-max <number>  write max in bytes
    --total-iops-sec-max <number>  total I/O operations max
    --read-iops-sec-max <number>  read I/O operations max
    --write-iops-sec-max <number>  write I/O operations max
    --size-iops-sec <number>  I/O size in bytes
    --config         affect next boot
    --live           affect running domain
    --current        affect current domain
```

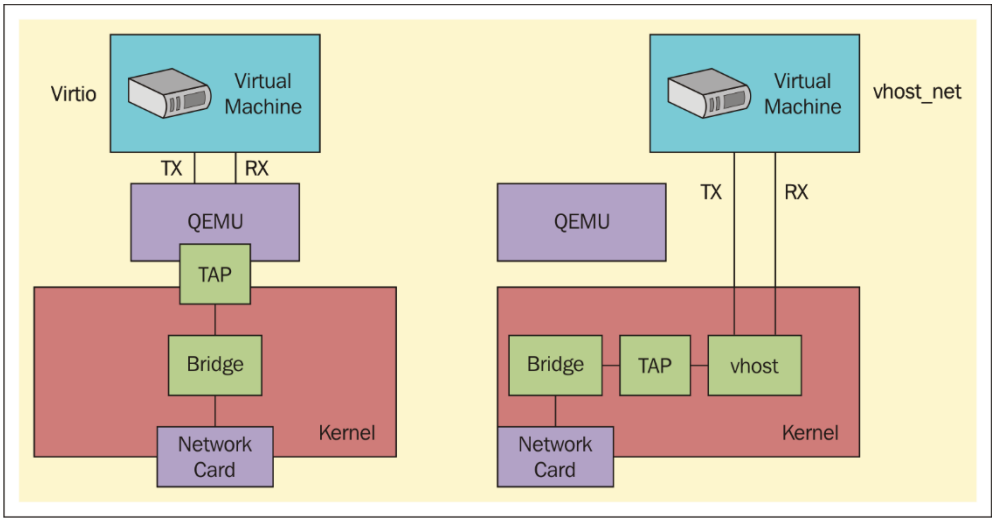


```
[humble-lap ]$ lsmod |grep vhost
vhost_net              20480  1
vhost                  32768  1 vhost_net
macvtap                20480  1 vhost_net
tun                    28672  4 vhost_net
[humble-lap ]$ modinfo vhost_net
filename:       /lib/modules/4.2.3-200.fc22.x86_64/kernel/drivers/vhost/vhost_net.ko.xz
alias:          devname:vhost-net
alias:          char-major-10-238
description:    Host kernel accelerator for virtio net
author:         Michael S. Tsirkin
license:        GPL v2
version:        0.0.1
srcversion:     1EDC0A4AEC45D8F033A71FE
depends:        vhost,tun,macvtap
intree:         Y
vermagic:       4.2.3-200.fc22.x86_64 SMP mod_unload
signer:         Fedora kernel signing key
sig_key:        6B:32:69:BB:F8:47:97:01:C8:03:15:FB:5F:36:8A:F9:24:52:07:BE
sig_hashalgo:   sha256
parm:           experimental_zcopytx:Enable Zero Copy TX; 1 -Enable; 0 - Disable (int)
[humble-lap ]$
[humble-lap ]$ modinfo --parameters vhost_net
experimental_zcopytx:Enable Zero Copy TX; 1 -Enable; 0 - Disable (int)
[humble-lap ]$
```

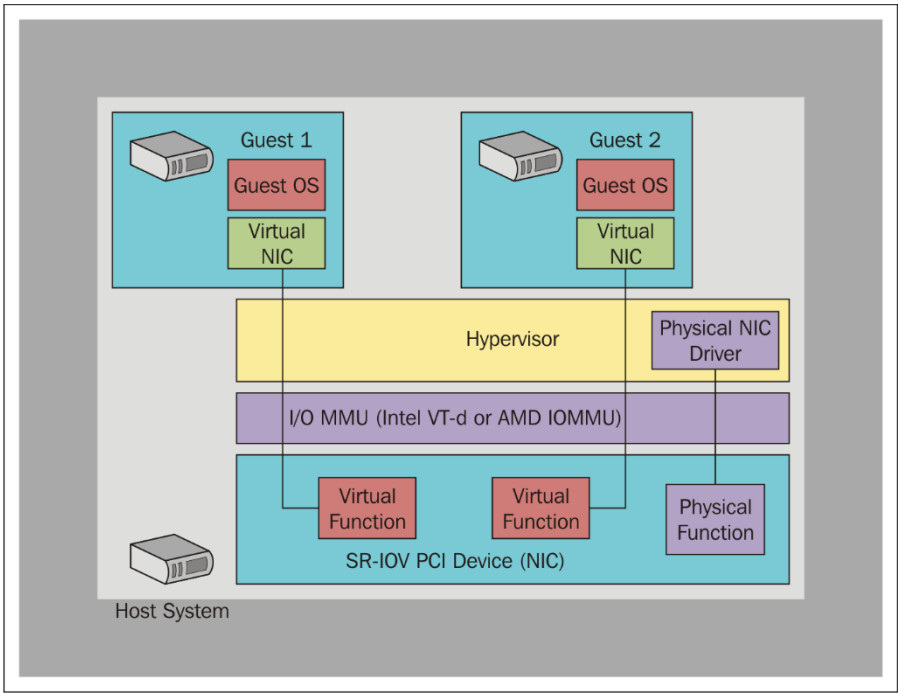
Guest 1
Guest OS
Virtual NIC
Guest 2
Guest OS
Virtual NIC
Hypervisor
Physical NIC Driver
I/O MMU (Intel VT-d or AMD IOMMU)
Virtual Function
Virtual Function
Physical Function
SR-IOV PCI Device (NIC)
Host System

Chapter 14: V2V and P2V Migration Tools

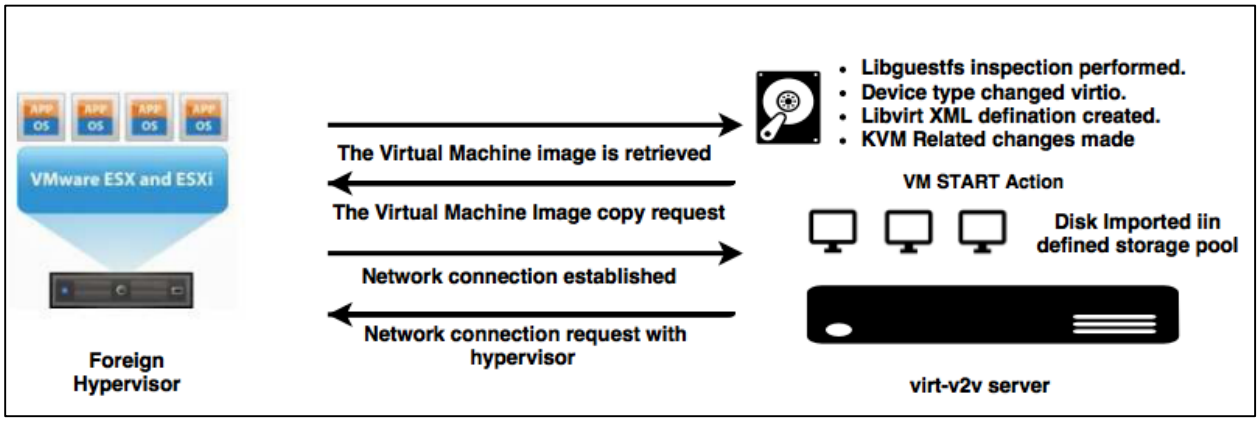

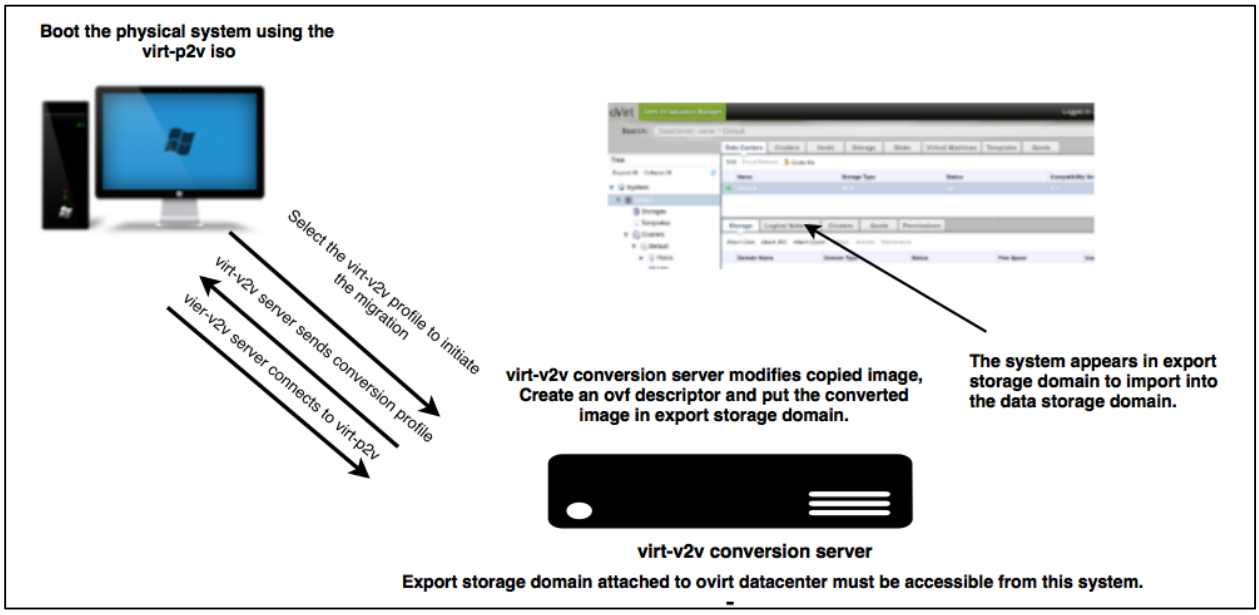

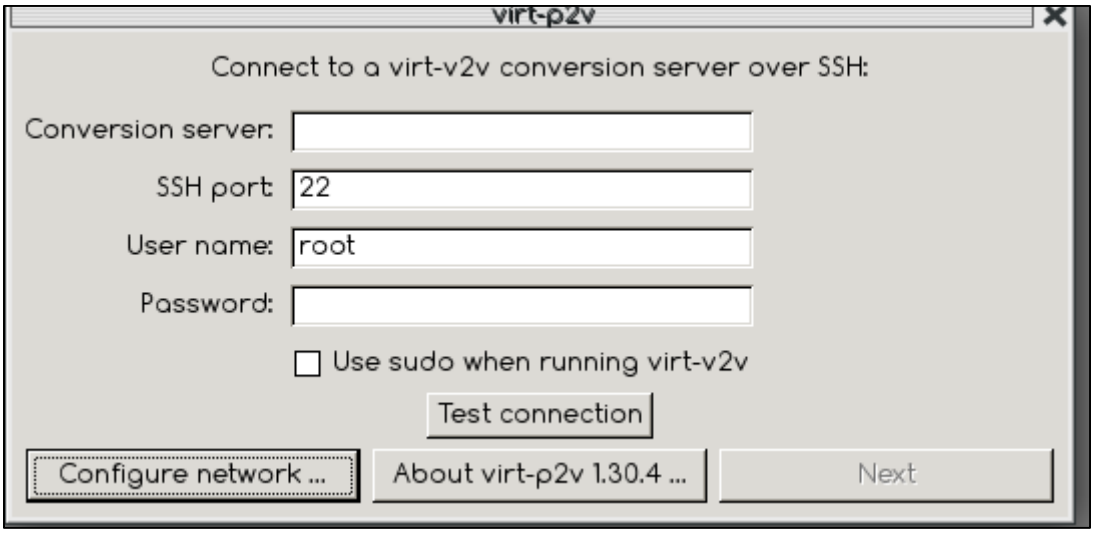

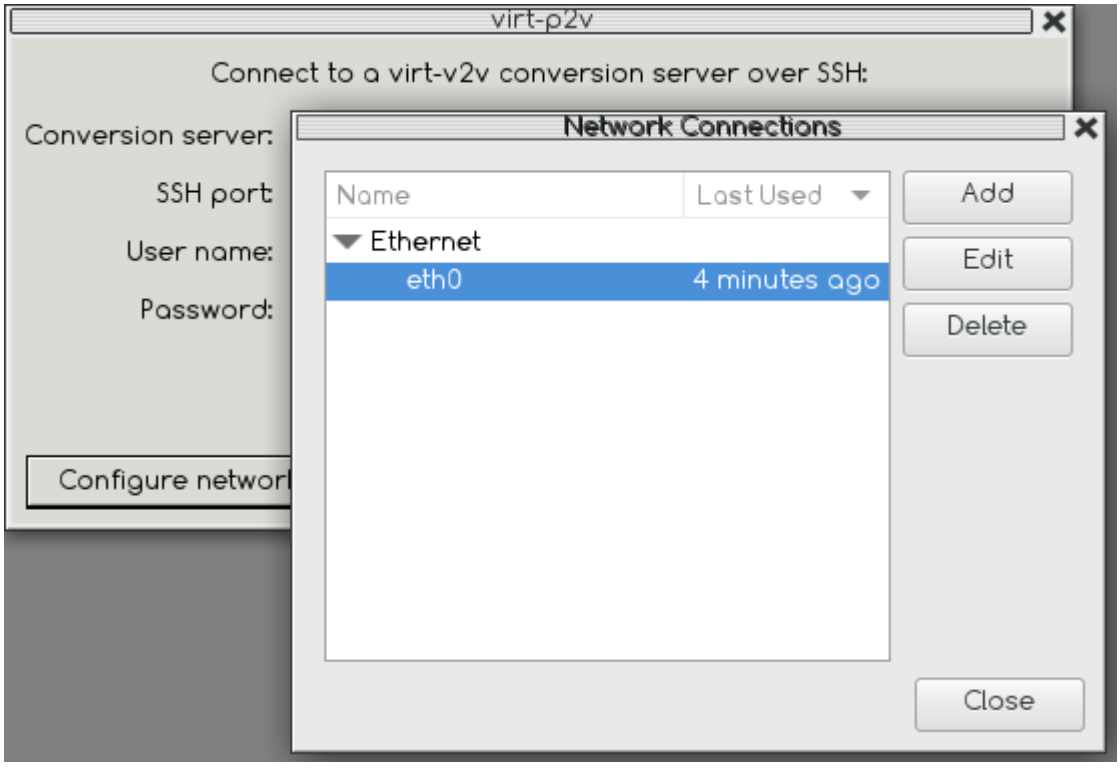
virt-p2v
Connect to a virt-v2v conversion server over SSH:
Conversion server:
SSH port:
User name:
Password:
Configure networ
Network Connections
Name
Last Used
Add
Ethernet
Edit
eth0
4 minutes ago
Delete
Close

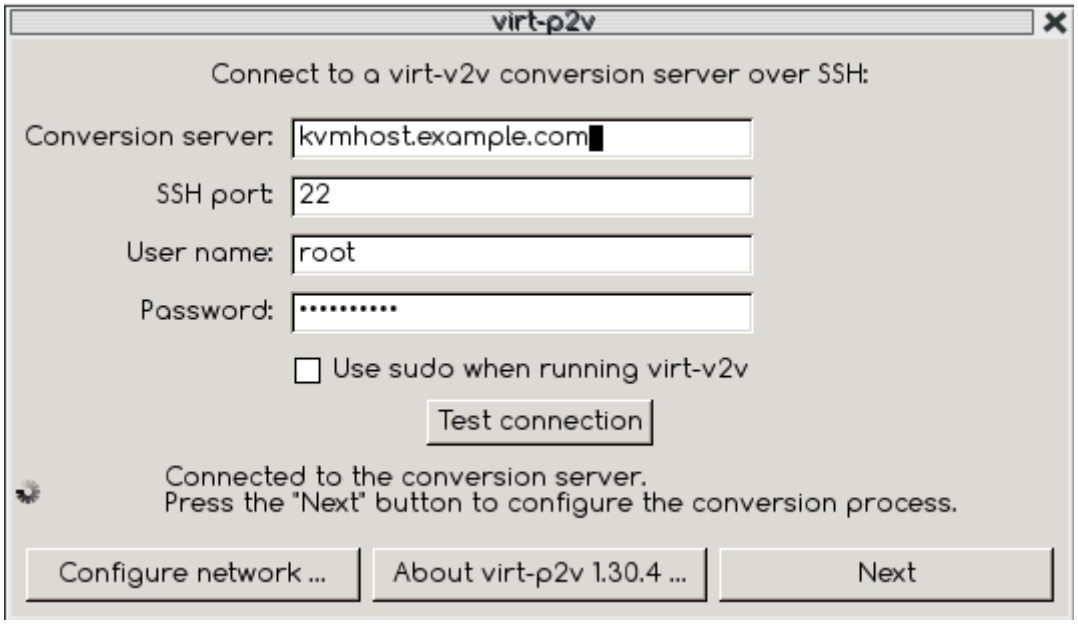
virt-p2v
Connect to a virt-v2v conversion server over SSH:
Conversion server:
kvmhost.example.com
SSH port:
22
User name:
root
Password:
Use sudo when running virt-v2v
Test connection
Connected to the conversion server.
Press the "Next" button to configure the conversion process.
Configure network ...
About virt-p2v 1.30.4 ...
Next

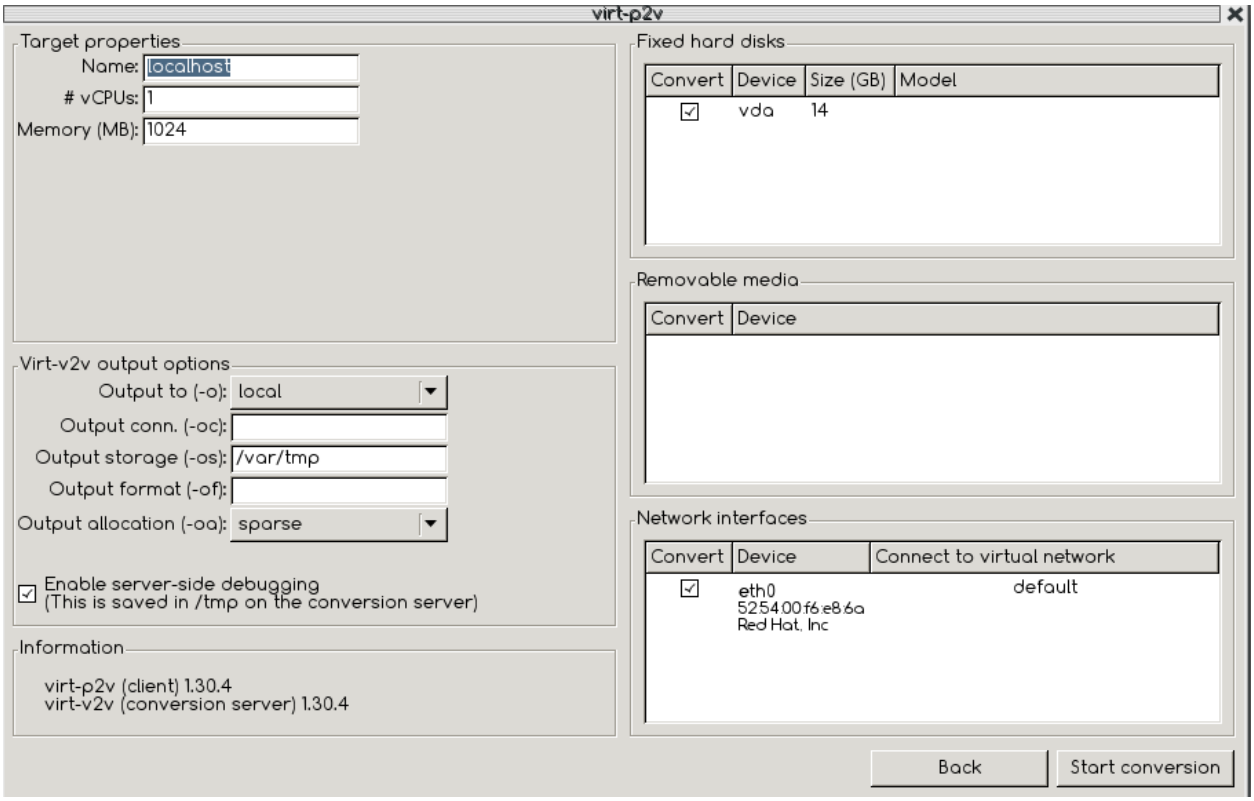
virt-p2v
Target properties
Name: localhost
vCPUs: 1
Memory (MB): 1024
Fixed hard disks
Convert
Device
Size (GB)
Model
vda
14
Removable media
Convert
Device
Virt-v2v output options
Output to (-o): local
Output conn. (-oc):
Output storage (-os): /var/tmp
Output format (-of):
Output allocation (-oa): sparse
Enable server-side debugging
(This is saved in /tmp on the conversion server)
Network interfaces
Convert
Device
Connect to virtual network
eth0
5254:00:f6:e8:6a
Red Hat, Inc
default
Information
virt-p2v (client) 1.30.4
virt-v2v (conversion server) 1.30.4
Back
Start conversion

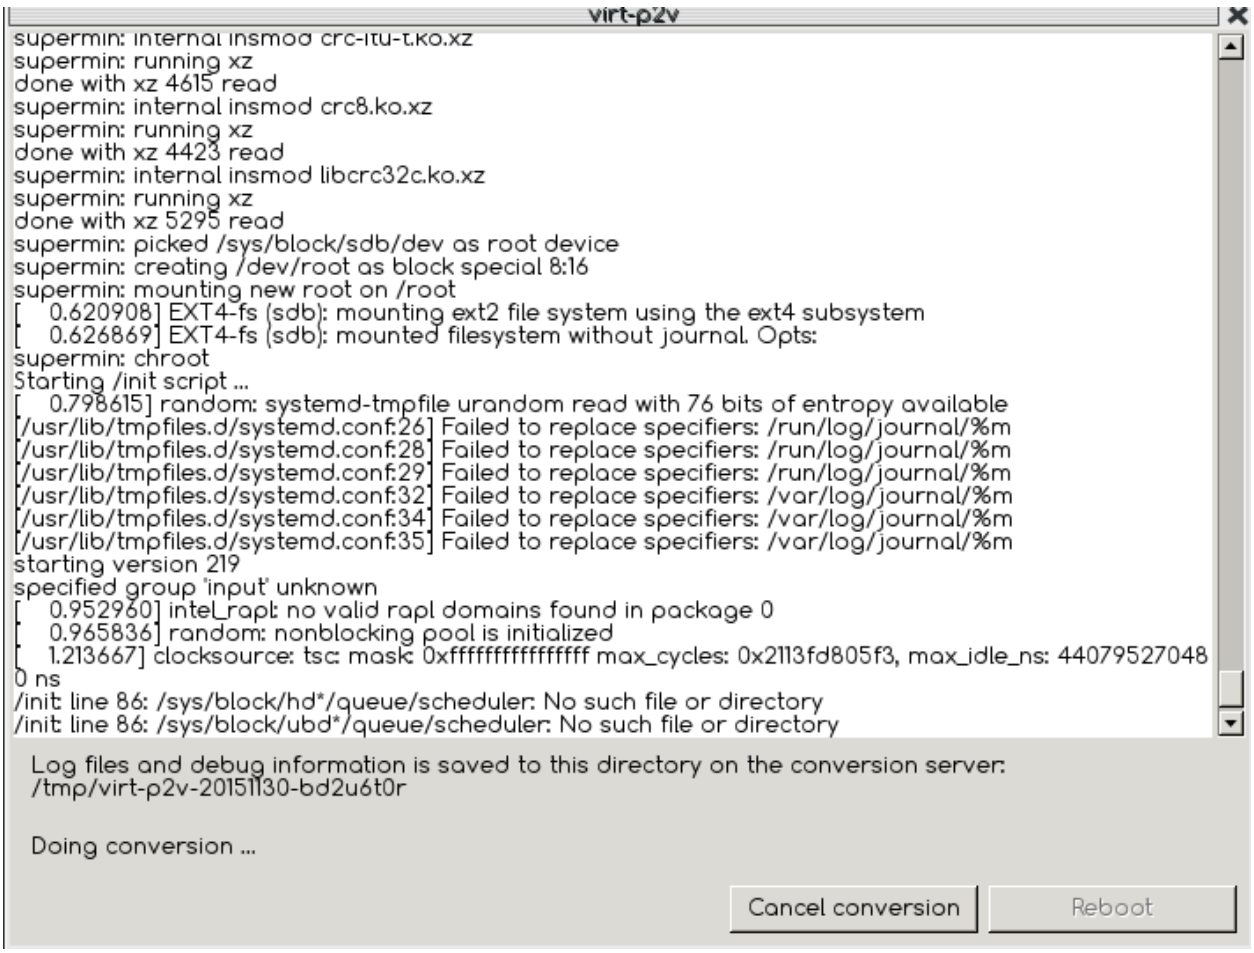
virt-p2v
supermin: internal insmod crc-itu-t.ko.xz
supermin: running xz
done with xz 4615 read
supermin: internal insmod crc8.ko.xz
supermin: running xz
done with xz 4423 read
supermin: internal insmod libcrc32c.ko.xz
supermin: running xz
done with xz 5295 read
supermin: picked /sys/block/sdb/dev as root device
supermin: creating /dev/root as block special 8:16
supermin: mounting new root on /root
[0.620908] EXT4-fs (sdb): mounting ext2 file system using the ext4 subsystem
[0.626869] EXT4-fs (sdb): mounted filesystem without journal. Opts:
supermin: chroot
Starting /init script ...
[0.798615] random: systemd-tmpfile urandom read with 76 bits of entropy available
[/usr/lib/tmpfiles.d/systemd.conf:26] Failed to replace specifiers: /run/log/journal/%m
[/usr/lib/tmpfiles.d/systemd.conf:28] Failed to replace specifiers: /run/log/journal/%m
[/usr/lib/tmpfiles.d/systemd.conf:29] Failed to replace specifiers: /run/log/journal/%m
[/usr/lib/tmpfiles.d/systemd.conf:32] Failed to replace specifiers: /var/log/journal/%m
[/usr/lib/tmpfiles.d/systemd.conf:34] Failed to replace specifiers: /var/log/journal/%m
[/usr/lib/tmpfiles.d/systemd.conf:35] Failed to replace specifiers: /var/log/journal/%m
starting version 219
specified group 'input' unknown
[0.952960] intel_rapl: no valid rapl domains found in package 0
[0.965836] random: nonblocking pool is initialized
[1.213667] clocksource: tsc: mask: 0xffffffffffffffff max_cycles: 0x2113fd805f3, max_idle_ns: 44079527048
0 ns
/init: line 86: /sys/block/hd*/queue/scheduler: No such file or directory
/init: line 86: /sys/block/ubd*/queue/scheduler: No such file or directory
Log files and debug information is saved to this directory on the conversion server:
/tmp/virt-p2v-20151130-bd2u6t0r
Doing conversion ...
Cancel conversion
Reboot

Appendix: Converting a Virtual Machine into a Hypervisor

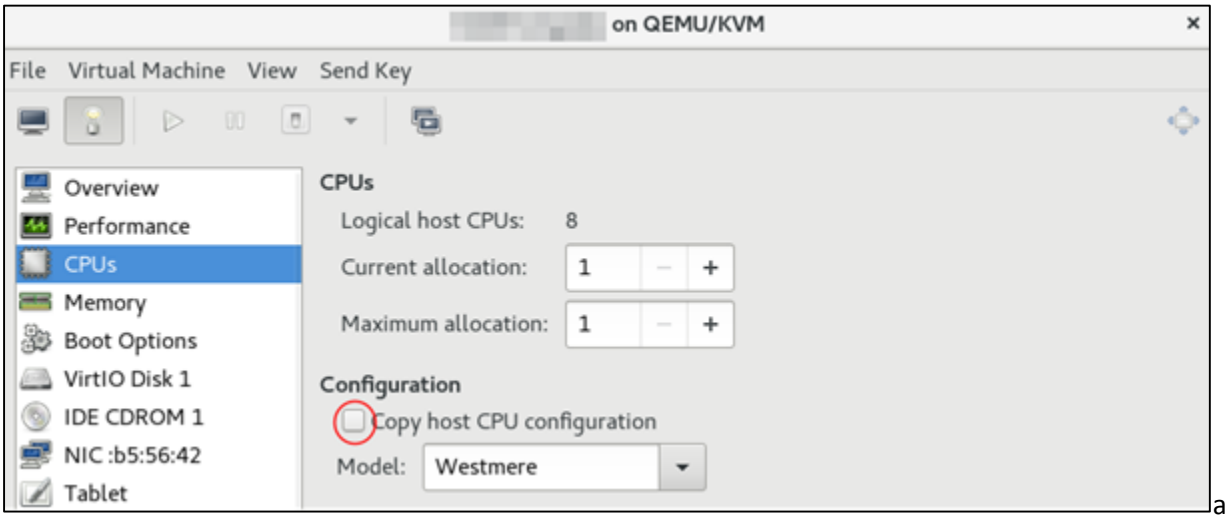

a