



Virtualisation with the Kernel-based Virtual Machine (KVM)

Version 2.2

Eng Diarmuid O'Briain, CEng, CISSP



Last updated: 02 May 2017

Copyright © 2017 Diarmuid Ó Briain

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the section entitled "GNU Free Documentation License".

Author

Diarmuid is a Chartered Engineer (CEng) with over 25 years experience in Telecommunications, Information Networking and Security. He has designed and implemented next-generation networks and information security solutions for major multi-national communications companies as well as serving as Chief Technical Officer for an Irish Internet Service Provider for over 5 years. Since 1999 he has also lectured on Telecommunications and Computing programmes at both the Dublin Institute of Technology (DIT), the Institute of Technology, Carlow (ITC) in Ireland and the College of Engineering, Design, Art and Technology (CEDAT) at Makerere University in Uganda.

Table of Contents

1. INTRODUCTION TO VIRTUALISATION.....	5
1.1 HYPERVERISOR.....	5
1.2 THIS DOCUMENT DEFAULTS.....	6
1.3 TERMINOLOGY.....	7
2. INTRODUCTION TO KVM.....	9
3. PRE-INSTALLATION CHECK.....	11
3.1 ENABLE VIRTUALISATION SUPPORT IN BIOS.....	11
4. INSTALLATION.....	13
4.1 INSTALL KVM PACKAGES.....	13
4.2 GNU/LINUX BRIDGE UTILITIES.....	13
4.3 CHECK INSTALLED VERSION.....	13
5. POST INSTALLATION SETUP.....	15
5.1 ETHERNET INTERFACE.....	15
5.2 ADD USERS TO GROUPS.....	16
5.3 VERIFY INSTALLATION.....	16
5.4 DISPLAY NODE INFORMATION.....	17
6. BUILDING VIRTUAL MACHINES.....	19
6.1 CREATION OF DOMAINS.....	19
6.2 DOMAIN DIRECTORY.....	19
6.3 OS VARIANT.....	19
6.4 BUILD FIRST DOMAIN.....	23
6.5 CONNECT TO A DOMAIN.....	24
6.6 BUILD A MICROSOFT WINDOWS DOMAIN.....	25
6.7 BUILD A DOMAIN OVER THE INTERNET.....	27
7. CONNECTING TO VIRTUAL MACHINES.....	29
7.1 CONNECT TO A REMOTE DOMAIN.....	29
8. MANAGE AND OPERATE THE DOMAINS.....	31
8.1 CHECK DOMAIN LIST.....	31
8.2 SHUTDOWN A DOMAIN.....	32
8.3 REBOOT A DOMAIN.....	32
8.4 TERMINATE A DOMAIN.....	32
8.5 START A DOMAIN.....	32
8.6 SUSPEND A DOMAIN.....	32
8.7 SAVE A DOMAIN.....	33
8.8 CLONING A DOMAIN.....	34
8.9 DELETE A DOMAIN.....	35
8.10 DOMAIN XML FILE.....	35
8.11 CREATING A SNAPSHOT.....	37
8.12 REVERT TO A PREVIOUS SNAPSHOT.....	39
8.13 DELETE A SNAPSHOT.....	41
9. VIRTUAL NETWORKS.....	43
9.1 CREATE A NEW NETWORK.....	44
9.2 CONFIGURE DOMAIN VMs TO CONNECT TO THE NEW NETWORK.....	45
9.3 CONFIGURE THE ROUTER VM01-DEBIAN_8.....	48
9.4 REVIEW THE BRIDGE CONTROL ON THE HOST.....	52

9.5 CONFIRM THAT THE GUESTS HAVE INTERNET ACCESS.....	52
10. GRAPHICAL MANAGEMENT.....	53
11. IMPORTING IMAGE FROM OTHER VIRTUALISATION PLATFORM.....	57
11.1 INSTALL THE KVM DOMAIN WITH THE IMPORT FILE.....	57
12. APPENDIX 1.....	59
12.1 CHECK IF KVM IS SUPPORTED.....	59
12.2 INSTALL THE KVM ELEMENTS.....	59
12.3 ADD USER TO THE LIBVIRT DAEMON GROUP.....	59
12.4 VM NETWORK.....	59
12.5 INSTALL A VM.....	60
12.6 ACCESSING THE VM's.....	63
12.7 IP ADDRESS ON KVM DOMAINS.....	64
13. ABBREVIATIONS.....	67
14. BIBLIOGRAPHY.....	69

Illustration Index

Illustration 1: Type-1 Bare-metal Hypervisor.....	5
Illustration 2: Type-2 Hypervisor.....	6
Illustration 3: KVM virtualisation block diagram.....	9
Illustration 4: Debian install.....	24
Illustration 5: Debian guest.....	25
Illustration 6: Windows guest install.....	26
Illustration 7: Windows guest.....	27
Illustration 8: Debian guest through SPICE viewer.....	29
Illustration 9: Windows guest through VNC viewer.....	30
Illustration 10: Virtual network.....	43
Illustration 11: Windows IPv4 address assignment.....	47
Illustration 12: Hypervisor 'Add connection' screen.....	53
Illustration 13: Domain dashboard view.....	54
Illustration 14: Domain details.....	55
Illustration 15: Create VM.....	60
Illustration 16: Select image.....	61
Illustration 17: Memory and CPU.....	61
Illustration 18: Create disk image.....	62
Illustration 19: Confirm configuration and finish.....	62
Illustration 20: Ubuntu Desktop VM.....	63
Illustration 21: Ubuntu Server VM.....	64

1. Introduction to Virtualisation

1.1 Hypervisor

A hypervisor or Virtual Machine Manager, is software that permits multiple OSs to share a single hardware host computer. Each OS appears to have the host's processor, memory, and other resources all to itself. However, the hypervisor is actually controlling the host processor and resources, allocating what is needed to each OS in turn and making sure that the guest Operating System (OS) on Virtual Machines (VM) cannot disrupt each other. The Kernel-based Virtual Machine (KVM) is the default hypervisor on most GNU's not Unix (GNU)/Linux distributions today. Before progressing to consider installing and operating it here are a few terms that are essential to both understand these examples and the technology.

1.1.1 Type-1 Hypervisor

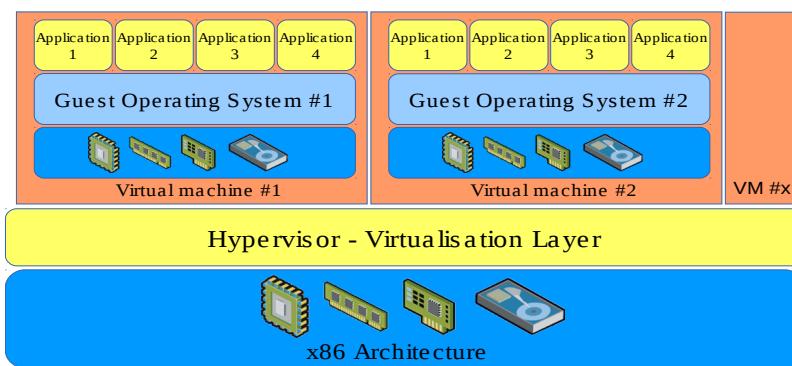


Illustration 1: Type-1 Bare-metal Hypervisor

A Type-1 hypervisor (baremetal hypervisor) is a client hypervisor that interacts directly with hardware that is being virtualised. It is completely independent of the host OS, unlike a Type-2 hypervisor, and boots before the host OS. Currently, Type-1 hypervisors are being used by all the major players in the desktop virtualisation space, including but not limited to the KVM, Xen hypervisor, VMware vSphere (ESX) hypervisor, and Microsoft Hyper-V Server.

1.1.2 Type-2 Hypervisor

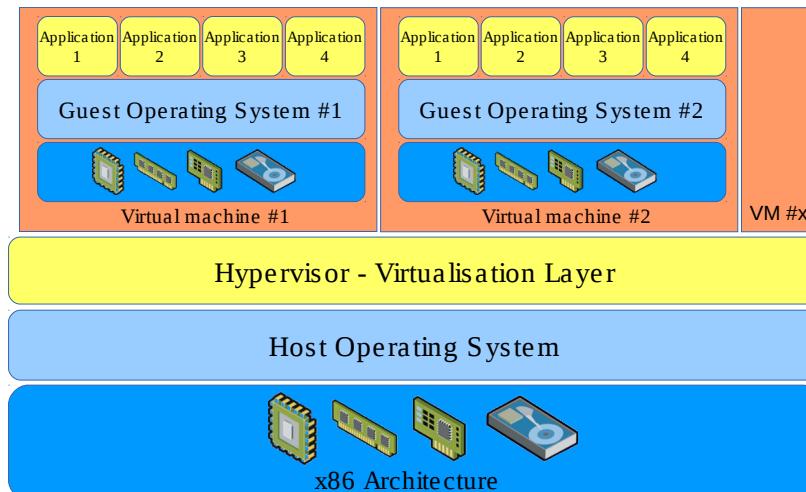


Illustration 2: Type-2 Hypervisor

A Type-2 hypervisor is a client hypervisor that sits on top of a host OS. Unlike a Type-1 hypervisor, a Type-2 hypervisor relies heavily on the host OS. It cannot boot until the host OS is already up and running and, if for any reason the host OS crashes, all end-users are affected. This is a big drawback of Type-2 hypervisors, as they are only as secure as the host OS on which they rely. Also, since Type-2 hypervisors depend on an OS, they are not in full control of the end user's machine. One advantage of a Type-2 hypervisor is that its functionality is practically identical on all of the host platforms, typically common file and image formats. All this results in the ability to run VMs created on one host on another host with a different host OS. i.e. you may have created a Microsoft Windows VM on GNU/Linux move it to another computer with a different host OS, say BSD UNIX and the hypervisor will be able to run the guest OS in a VM. Some examples include Oracle VirtualBox, VMware Server and Microsoft Virtual Server.

1.1.3 Type-1 -vs- Type-2 Hypervisor

For high performance, Return on Investment (ROI) and scalability a Type-1 hypervisor is far superior to a Type-2. This is due to the independence of the Type-1 hypervisor from the operating system.

1.2 This document defaults

The following locations have been used for files and this programme is ably assisted by the famous programmer Ada Lovelace.

- **/virt/kvm/images** - Directory to store VMs.
- **/virt/iso** - Directory to store ISO images.
- User: **Ada Lovelace** Username: **alovelace** Password: **babbage**
- Hypervisor: **Debian 8.2 GNU/Linux**

1.3 Terminology

- **node** is a single physical machine.
- **hypervisor** is a layer of software allowing to virtualise a node in a set of VMs with possibly different configurations than the node itself.
- **KVM** is the GNU/Linux Hypervisor.
- **domain** is an instance of an OS (or subsystem in the case of container virtualisation) running on a virtualised machine provided by the hypervisor.
- **QEMU** is a Quick EMULATOR (QEMU), a generic and open source machine emulator and virtualiser for Input/Output (I/O) hardware emulation.
- **libvirt** is a toolkit virtualisation management system to interact with the virtualisation capabilities of recent versions of Linux, KVM, Xen and LinuX Container (LXC).
- **libvirtd** is the server side daemon component of the libvirt virtualisation management system.
- **virsh** is a command line interface tool for managing guests and the hypervisor.
- **virt-manager** is a desktop user interface for managing virtual machines through libvirt.
- **virt-viewer** is a lightweight interface for interacting with the graphical display of virtualised guest OS. It can display Virtual Network Computing (VNC) or Simple Protocol for Independent Computing Environments (SPICE) and uses libvirt to lookup the graphical connection details.
- **virt-clone** is a command line tool for cloning existing inactive guests. It copies the disk images, and defines a configuration with new name, Universally Unique Identifier (UUID) and Medium Access Control (MAC) address pointing to the copied disks.
- **spicec** is a SPICE graphical client. SPICE is said to be faster than VNC as a remote desktop protocol.
- **vncviewer** is a VNC graphical client.

This page is intentionally blank

2. Introduction to KVM

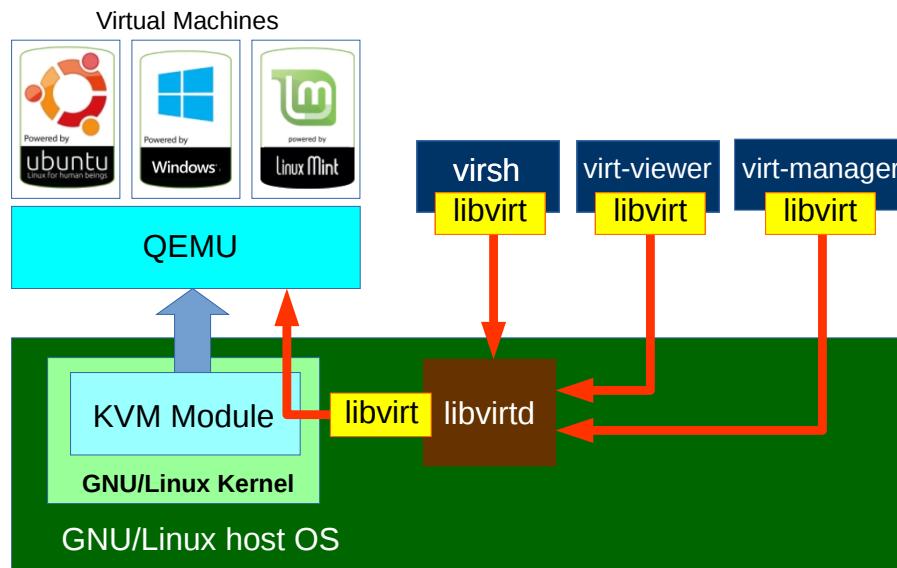


Illustration 3: KVM virtualisation block diagram

The KVM hypervisor is included in major GNU/Linux releases as it has become the hypervisor of choice within the GNU/Linux community so it is available within the various distribution repositories. In actual fact KVM is a GNU/Linux kernel module that permits programs within user space to access either the Intel or AMD processor virtualisation features. As a result KVM Virtual Machines (VM) actually run as user space processes. KVM uses the QEMU, a generic and open source machine emulator and virtualiser for I/O hardware emulation. It can emulate a variety of processors on a guest processor and combined with the KVM kernel module it can approach native speeds. All combinations of 32-bit and 64-bit host and guest systems are supported, except 64-bit guests on 32-bit hosts.

KVM is managed via the libvirt API and tools such as *virsh*, *virtinstall*, *virt-clone*, *virt-viewer* and *virt-manager*.

KVM is a Type-1 hypervisor that runs directly on x86 hardware. The GNU/Linux interface makes it look like it is a hosted hypervisor running on it, but in fact each VM is running on the bare metal with the host GNU/Linux OS providing a launchpad for the hypervisor and then engaging in a co-processing relationship with the hypervisor.

On x86 hardware, KVM relies on the hardware virtualisation instructions that are embedded in the processors and therefore these advanced chipset features must be enabled. Using these instructions the hypervisor and each guest VM run directly on the bare metal, and most of the resource translations are performed by the hardware.

This page is intentionally blank

3. Pre-installation check

3.1 Enable virtualisation support in BIOS

To support Hardware-assisted Virtual Machine (HVM) guests, virtualisation extensions need to be enabled in the Basic Input/Output System (BIOS). In the BIOS the Virtualise option appears under “Advanced Chipset Features” as one of the following:

- Enable Virtualisation Technology for x86 architectures (VT-x).
- Enable Intel VT.
- Vanderpool Technology.

Also enable:

- Intel Virtualisation Technology for Directed I/O (VT-d)

Confirm that the hardware virtualisation is now supported by the CPU by searching for *Virtual Machine eXtensions* (VMX) to see if the computer has an Intel processor or *Secure Virtual Machine* (SVM) for AMD support if the hardware has an AMD processor.

Check that the CPU supports hardware virtualisation. 0 means that the CPU doesn't support hardware virtualisation while > 0 means it does but it still needs to be enabled in the BIOS. (ie. *vmx* or *svm* has appeared x number of times in the output of the command.

```
ada:~$ egrep -c '(vmx|svm)' /proc/cpuinfo  
4
```

Check if a 64 bit kernel is running. 0 means that the Central Processing Unit (CPU) is not 64-bit. *Long Mode* (LM) equates to a 64-bit CPU.

```
ada:~$ egrep -c ' lm ' /proc/cpuinfo  
8  
  
ada:~$ uname -m  
x86_64
```

This page is intentionally blank

4. Installation

KVM requires a number of elements to operate.

4.1 Install KVM packages

- **libvirt** is a C toolkit to interact with the virtualisation capabilities of GNU/Linux. The library provides a C API for different virtualisation mechanisms and currently supports QEMU, KVM, Xen, OpenVZ, LXC, and VirtualBox.
- **qemu-kvm** permits the running of multiple virtual computers, each running unmodified GNU/Linux or Windows images on X86 hardware. Each virtual machine has private virtualised hardware: a network card, disk, graphics adapter, etc.
- **virt-manager** - graphical user interface

```
ada:~$ sudo apt-get install qemu-kvm libvirt-bin virt-manager  
libosinfo-bin
```

4.2 GNU/Linux Bridge utilities

Without a bridge KVM/QEMU VMs will only have network access to other VMs on the same server and to the host itself via a shared private network 192.168.122.0. To allow VMs access to the LAN, create a network bridge on the host.

```
ada:~$ sudo apt-get install bridge-utils
```

4.3 Check installed version

The KVM version follows the kernel version so a quick way to check is to see the kernel version of the GNU/Linux operating system.

```
ada:~$ uname -r  
4.10.0-20-generic
```

The KVM version command gives the QEMU version as that is the core element of KVM. Another element of the system worth noting is the virsh management user interface version.

```
ada:~$ kvm --version  
QEMU emulator version 2.8.0 (Debian 1:2.8+dfsg-3ubuntu2.1)  
Copyright (c) 2003-2016 Fabrice Bellard and the QEMU Project  
developers
```

```
ada:~$ virsh --version  
2.5.0
```

This page is intentionally blank

5. Post installation setup

5.1 Ethernet interface

Edit the /etc/network/interfaces file by creating a bridge, put *eth0* into the bridge and transfer the Internet Protocol (IP) addressing information to the bridge. Restart networking.

```
ada:~$ sudo vi /etc/network/interfaces
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

source /etc/network/interfaces.*

# The loopback network interface
auto lo
iface lo inet loopback

# The primary interface eth0
auto eth0
iface eth0 inet manual

auto br0
iface br0 inet static
    address 213.74.34.100
    netmask 255.255.255.0
    network 213.74.34.0
    broadcast 213.74.34.255
    gateway 213.74.34.1
    dns-nameservers 8.8.8.8
    bridge_ports eth0
    bridge_fd 9
    bridge_hello 2
    bridge_maxage 12
    bridge_stp off

:wq!
```

Here is an explanation on some of the bridge settings.

- **bridge_ports eth0** - Define ports within the bridge.
- **bridge_fd 9** - Set bridge forward delay to 9 seconds (Default: 15).
- **bridge_maxage 12** - Set max message age to 12 seconds (Default: 20).
- **bridge_stp off** - Disables Spanning Tree Protocol.

Restart the *networking* service.

```
ada:~$ sudo /etc/init.d/networking restart
[ ok ] Restarting networking (via systemctl): networking.service.
```

Now confirm the IP settings.

```
ada:~$ ip addr list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast master br0 state UP group default qlen 1000
    link/ether 30:5a:3a:08:39:21 brd ff:ff:ff:ff:ff:ff
    inet 213.74.34.100/24 brd 213.74.34.255 scope global eth0
        valid_lft forever preferred_lft forever
3: br0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group default
    link/ether 30:5a:3a:08:39:21 brd ff:ff:ff:ff:ff:ff
    inet 213.74.34.100/24 brd 213.74.34.255 scope global br0
        valid_lft forever preferred_lft forever
    inet6 fe80::325a:3aff:fe08:3921/64 scope link
        valid_lft forever preferred_lft forever
```

5.2 Add Users to Groups

Add the user to the *libvird* and *kvm* groups. Then logout and back in to activate.

```
ada:~$ sudo adduser `id -un` libvirt
Adding user 'alovelace' to group 'libvirt' ...

ada:~$ sudo adduser `id -un` kvm
Adding user 'alovelace' to group 'kvm' ...
```

Confirm entries.

```
ada:~$ sudo egrep '(kvm|libvirt)' /etc/group
libvirt:x:124:alovelace
kvm:x:125:alovelace
libvirt-qemu:x:126:libvirt-qemu
```

Logout of shell and log back in to join the groups. Use the *id* command to verify.

```
ada:~$ id
uid=1000(alovelace) gid=1000(alovelace)
groups=1000(alovelace),24(cdrom),25(floppy),27(sudo),29(audio),
30(dip),44(video),46(plugdev),108(netdev),111(scanner),
115(bluetooth),124(libvirt),125(kvm)
```

5.3 Verify installation

If the installation is correct then a list like this will be presented. As no VMs are generated as yet this list will be empty.

```
ada:~$ virsh -c qemu:///system

virsh # list
      Id      Name           State
```

5.4 Display node information

To display node information for the hypervisor, the host machine that supports the virtualisation process use the *nodeinfo* command.

```
virsh # nodeinfo
CPU model:          x86_64
CPU(s):             4
CPU frequency:      800 MHz
CPU socket(s):      1
Core(s) per socket: 4
Thread(s) per core: 1
NUMA cell(s):       1
Memory size:        16151212 KiB
```

This page is intentionally blank

6. Building Virtual Machines

6.1 Creation of domains

The tools that will be used can be seen by typing `virt` and hitting <TAB>.

```
ada:~$ virt
virt-clone  virt-host-validate  virt-login-shell  virt-viewer
virt-convert  virt-image  virt-manager  virt-xml
virtfs-proxy-helper  virt-install  virt-pki-validate
virt-xml-validate
```

6.2 Domain directory

By default images are stored in the directory `/var/lib/libvirt/images/` which was created automatically when `qemu-kvm` was installed. However the machine in use has a separate harddrive `/dev/sdb` for VMs and it is mounted as `/virt`, therefore images will be stored in `/virt/kvm/images/`.

6.3 OS variant

It is useful to optimise the guest configuration for the specific OS to be installed. Get a list of OS variants as follows.

Short ID	Name	Version	ID
altlinux1.0	Mandrake RE Spring 2001	1.0	http://altlinux.org/altlinux/1.0
altlinux2.0	ALT Linux 2.0	2.0	http://altlinux.org/altlinux/2.0
altlinux2.2	ALT Linux 2.2	2.2	http://altlinux.org/altlinux/2.2
altlinux2.4	ALT Linux 2.4	2.4	http://altlinux.org/altlinux/2.4
altlinux3.0	ALT Linux 3.0	3.0	http://altlinux.org/altlinux/3.0
altlinux4.0	ALT Linux 4.0	4.0	http://altlinux.org/altlinux/4.0
altlinux4.1	ALT Linux 4.1	4.1	http://altlinux.org/altlinux/4.1
altlinux5.0	ALT Linux 5.0	5.0	http://altlinux.org/altlinux/5.0
altlinux6.0	ALT Linux 6.0	6.0	http://altlinux.org/altlinux/6.0
altlinux7.0	ALT Linux 7.0	7.0	http://altlinux.org/altlinux/7.0
centos6.0	CentOS 6.0	6.0	http://centos.org/centos/6.0
centos6.1	CentOS 6.1	6.1	http://centos.org/centos/6.1
centos6.2	CentOS 6.2	6.2	http://centos.org/centos/6.2
centos6.3	CentOS 6.3	6.3	http://centos.org/centos/6.3
centos6.4	CentOS 6.4	6.4	http://centos.org/centos/6.4
centos6.5	CentOS 6.5	6.5	http://centos.org/centos/6.5
centos6.6	CentOS 6.6	6.6	http://centos.org/centos/6.6
centos6.7	CentOS 6.7	6.7	http://centos.org/centos/6.7
centos7.0	CentOS 7.0	7.0	http://centos.org/centos/7.0
debian1.1	Debian Buzz	1.1	http://debian.org/debian/1.1
debian1.2	Debian Rex	1.2	http://debian.org/debian/1.2
debian1.3	Debian Bo	1.3	http://debian.org/debian/1.3
debian2.0	Debian Hamm	2.0	http://debian.org/debian/2.0
debian2.1	Debian Slink	2.1	http://debian.org/debian/2.1
debian2.2	Debian Potato	2.2	http://debian.org/debian/2.2
debian3	Debian Woody	3	http://debian.org/debian/3
debian3.1	Debian Sarge	3.1	http://debian.org/debian/3.1
debian4	Debian Etch	4	http://debian.org/debian/4
debian5	Debian Lenny	5	http://debian.org/debian/5
debian6	Debian Squeeze	6	http://debian.org/debian/6
debian7	Debian Wheezy	7	http://debian.org/debian/7
debian8	Debian Jessie	8	http://debian.org/debian/8
debian9	Debian Stretch	9	http://debian.org/debian/9
debiantesting	Debian Testing	testing	http://debian.org/debian/testing
fedora-unknown	Fedora	unknown	http://fedoraproject.org/fedora/unknown
fedora1	Fedora Core 1	1	http://fedoraproject.org/fedora/1
fedora10	Fedora 10	10	http://fedoraproject.org/fedora/10
fedora11	Fedora 11	11	http://fedoraproject.org/fedora/11
fedora12	Fedora 12	12	http://fedoraproject.org/fedora/12
fedora13	Fedora 13	13	http://fedoraproject.org/fedora/13
fedora14	Fedora 14	14	http://fedoraproject.org/fedora/14
fedora15	Fedora 15	15	http://fedoraproject.org/fedora/15
fedora16	Fedora 16	16	http://fedoraproject.org/fedora/16
fedora17	Fedora 17	17	http://fedoraproject.org/fedora/17
fedora18	Fedora 18	18	http://fedoraproject.org/fedora/18
fedora19	Fedora 19	19	http://fedoraproject.org/fedora/19
fedora2	Fedora Core 2	2	http://fedoraproject.org/fedora/2
fedora20	Fedora 20	20	http://fedoraproject.org/fedora/20
fedora21	Fedora 21	21	http://fedoraproject.org/fedora/21
fedora22	Fedora 22	22	http://fedoraproject.org/fedora/22
fedora23	Fedora 23	23	http://fedoraproject.org/fedora/23
fedora24	Fedora 24	24	http://fedoraproject.org/fedora/24
fedora25	Fedora 25	25	http://fedoraproject.org/fedora/25
fedora3	Fedora Core 3	3	http://fedoraproject.org/fedora/3
fedora4	Fedora Core 4	4	http://fedoraproject.org/fedora/4
fedora5	Fedora Core 5	5	http://fedoraproject.org/fedora/5
fedora6	Fedora Core 6	6	http://fedoraproject.org/fedora/6
fedora7	Fedora 7	7	http://fedoraproject.org/fedora/7
fedora8	Fedora 8	8	http://fedoraproject.org/fedora/8

fedora9	Fedora 9	9	http://fedoraproject.org/fedora/9
freebsd1.0	FreeBSD 1.0	1.0	http://freebsd.org/freebsd/1.0
freebsd10.0	FreeBSD 10.0	10.0	http://freebsd.org/freebsd/10.0
freebsd10.1	FreeBSD 10.1	10.1	http://freebsd.org/freebsd/10.1
freebsd10.2	FreeBSD 10.2	10.2	http://freebsd.org/freebsd/10.2
freebsd10.3	FreeBSD 10.3	10.3	http://freebsd.org/freebsd/10.3
freebsd11.0	FreeBSD 11.0	11.0	http://freebsd.org/freebsd/11.0
freebsd2.0	FreeBSD 2.0	2.0	http://freebsd.org/freebsd/2.0
freebsd2.0.5	FreeBSD 2.0.5	2.0.5	http://freebsd.org/freebsd/2.0.5
freebsd2.2.8	FreeBSD 2.2.8	2.2.8	http://freebsd.org/freebsd/2.2.8
freebsd2.2.9	FreeBSD 2.2.9	2.2.9	http://freebsd.org/freebsd/2.2.9
freebsd3.0	FreeBSD 3.0	3.0	http://freebsd.org/freebsd/3.0
freebsd3.2	FreeBSD 3.2	3.2	http://freebsd.org/freebsd/3.2
freebsd4.0	FreeBSD 4.0	4.0	http://freebsd.org/freebsd/4.0
freebsd4.1	FreeBSD 4.1	4.1	http://freebsd.org/freebsd/4.1
freebsd4.10	FreeBSD 4.10	4.10	http://freebsd.org/freebsd/4.10
freebsd4.11	FreeBSD 4.11	4.11	http://freebsd.org/freebsd/4.11
freebsd4.2	FreeBSD 4.2	4.2	http://freebsd.org/freebsd/4.2
freebsd4.3	FreeBSD 4.3	4.3	http://freebsd.org/freebsd/4.3
freebsd4.4	FreeBSD 4.4	4.4	http://freebsd.org/freebsd/4.4
freebsd4.5	FreeBSD 4.5	4.5	http://freebsd.org/freebsd/4.5
freebsd4.6	FreeBSD 4.6	4.6	http://freebsd.org/freebsd/4.6
freebsd4.7	FreeBSD 4.7	4.7	http://freebsd.org/freebsd/4.7
freebsd4.8	FreeBSD 4.8	4.8	http://freebsd.org/freebsd/4.8
freebsd4.9	FreeBSD 4.9	4.9	http://freebsd.org/freebsd/4.9
freebsd5.0	FreeBSD 5.0	5.0	http://freebsd.org/freebsd/5.0
freebsd5.1	FreeBSD 5.1	5.1	http://freebsd.org/freebsd/5.1
freebsd5.2	FreeBSD 5.2	5.2	http://freebsd.org/freebsd/5.2
freebsd5.2.1	FreeBSD 5.2.1	5.2.1	http://freebsd.org/freebsd/5.2.1
freebsd5.3	FreeBSD 5.3	5.3	http://freebsd.org/freebsd/5.3
freebsd5.4	FreeBSD 5.4	5.4	http://freebsd.org/freebsd/5.4
freebsd5.5	FreeBSD 5.5	5.5	http://freebsd.org/freebsd/5.5
freebsd6.0	FreeBSD 6.0	6.0	http://freebsd.org/freebsd/6.0
freebsd6.1	FreeBSD 6.1	6.1	http://freebsd.org/freebsd/6.1
freebsd6.2	FreeBSD 6.2	6.2	http://freebsd.org/freebsd/6.2
freebsd6.3	FreeBSD 6.3	6.3	http://freebsd.org/freebsd/6.3
freebsd6.4	FreeBSD 6.4	6.4	http://freebsd.org/freebsd/6.4
freebsd7.0	FreeBSD 7.0	7.0	http://freebsd.org/freebsd/7.0
freebsd7.1	FreeBSD 7.1	7.1	http://freebsd.org/freebsd/7.1
freebsd7.2	FreeBSD 7.2	7.2	http://freebsd.org/freebsd/7.2
freebsd7.3	FreeBSD 7.3	7.3	http://freebsd.org/freebsd/7.3
freebsd7.4	FreeBSD 7.4	7.4	http://freebsd.org/freebsd/7.4
freebsd8.0	FreeBSD 8.0	8.0	http://freebsd.org/freebsd/8.0
freebsd8.1	FreeBSD 8.1	8.1	http://freebsd.org/freebsd/8.1
freebsd8.2	FreeBSD 8.2	8.2	http://freebsd.org/freebsd/8.2
freebsd8.3	FreeBSD 8.3	8.3	http://freebsd.org/freebsd/8.3
freebsd8.4	FreeBSD 8.4	8.4	http://freebsd.org/freebsd/8.4
freebsd9.0	FreeBSD 9.0	9.0	http://freebsd.org/freebsd/9.0
freebsd9.1	FreeBSD 9.1	9.1	http://freebsd.org/freebsd/9.1
freebsd9.2	FreeBSD 9.2	9.2	http://freebsd.org/freebsd/9.2
freebsd9.3	FreeBSD 9.3	9.3	http://freebsd.org/freebsd/9.3
gnome-continuous-3.10	GNOME 3.10	3.10	http://gnome.org/gnome-continuous/3.10
gnome-continuous-3.12	GNOME 3.12	3.12	http://gnome.org/gnome-continuous/3.12
gnome-continuous-3.14	GNOME 3.14	3.14	http://gnome.org/gnome-continuous/3.14
gnome3.6	GNOME 3.6	3.6	http://gnome.org/gnome/3.6
gnome3.8	GNOME 3.8	3.8	http://gnome.org/gnome/3.8
macosx10.0	MacOS X Cheetah	10.0	http://apple.com/macosx/10.0
macosx10.1	MacOS X Puma	10.1	http://apple.com/macosx/10.1
macosx10.2	MacOS X Jaguar	10.2	http://apple.com/macosx/10.2
macosx10.3	MacOS X Panther	10.3	http://apple.com/macosx/10.3
macosx10.4	MacOS X Tiger	10.4	http://apple.com/macosx/10.4
macosx10.5	MacOS X Leopard	10.5	http://apple.com/macosx/10.5
macosx10.6	MacOS X Snow Leopard	10.6	http://apple.com/macosx/10.6
macosx10.7	MacOS X Lion	10.7	http://apple.com/macosx/10.7
mageia1	Mageia 1	1	http://mageia.org/mageia/1
mageia2	Mageia 2	2	http://mageia.org/mageia/2
mageia3	Mageia 3	3	http://mageia.org/mageia/3
mandrake10.0	Mandrake Linux 10.0	10.0	http://mandriva.com/mandrake/10.0
mandrake10.1	Mandrake Linux 10.1	10.1	http://mandriva.com/mandrake/10.1
mandrake10.2	Mandrake Linux 10.2	10.2	http://mandriva.com/mandrake/10.2
mandrake5.1	Mandrake Linux 5.1	5.1	http://mandriva.com/mandrake/5.1
mandrake5.2	Mandrake Linux 5.2	5.2	http://mandriva.com/mandrake/5.2
mandrake5.3	Mandrake Linux 5.3	5.3	http://mandriva.com/mandrake/5.3
mandrake6.0	Mandrake Linux 6.0	6.0	http://mandriva.com/mandrake/6.0
mandrake6.1	Mandrake Linux 6.1	6.1	http://mandriva.com/mandrake/6.1
mandrake7.0	Mandrake Linux 7.0	7.0	http://mandriva.com/mandrake/7.0
mandrake7.1	Mandrake Linux 7.1	7.1	http://mandriva.com/mandrake/7.1
mandrake7.2	Mandrake Linux 7.2	7.2	http://mandriva.com/mandrake/7.2
mandrake8.0	Mandrake Linux 8.0	8.0	http://mandriva.com/mandrake/8.0
mandrake8.1	Mandrake Linux 8.1	8.1	http://mandriva.com/mandrake/8.1
mandrake8.2	Mandrake Linux 8.2	8.2	http://mandriva.com/mandrake/8.2
mandrake9.0	Mandrake Linux 9.0	9.0	http://mandriva.com/mandrake/9.0
mandrake9.1	Mandrake Linux 9.1	9.1	http://mandriva.com/mandrake/9.1
mandrake9.2	Mandrake Linux 9.2	9.2	http://mandriva.com/mandrake/9.2
mandriva2006.0	Mandriva Linux 2006.0	2006.0	http://mandriva.com/mandriva/2006.0
mandriva2007	Mandriva Linux 2007	2007	http://mandriva.com/mandriva/2007
mandriva2007.1	Mandriva Linux 2007 Spring	2007.1	http://mandriva.com/mandriva/2007.1
mandriva2008.0	Mandriva Linux 2008	2008.0	http://mandriva.com/mandriva/2008.0
mandriva2008.1	Mandriva Linux 2008 Spring	2008.1	http://mandriva.com/mandriva/2008.1
mandriva2009.0	Mandriva Linux 2009	2009.0	http://mandriva.com/mandriva/2009.0
mandriva2009.1	Mandriva Linux 2009 Spring	2009.1	http://mandriva.com/mandriva/2009.1
mandriva2010.0	Mandriva Linux 2010	2010.0	http://mandriva.com/mandriva/2010.0
mandriva2010.1	Mandriva Linux 2010 Spring	2010.1	http://mandriva.com/mandriva/2010.1
mandriva2010.2	Mandriva Linux 2010.2	2010.2	http://mandriva.com/mandriva/2010.2
mandriva2011	Mandriva Linux 2011	2011	http://mandriva.com/mandriva/2011
mbs1.0	Mandriva Business Server 1.0	1.0	http://mandriva.com/mbs/1.0
mes5	Mandriva Enterprise Server 5.0	5.0	http://mandriva.com/mes/5.0
mes5.1	Mandriva Enterprise Server 5.1	5.1	http://mandriva.com/mes/5.1
msdos6.22	Microsoft MS-DOS 6.22	6.22	http://microsoft.com/msdos/6.22
netbsd0.8	NetBSD 0.8	0.8	http://netbsd.org/netbsd/0.8
netbsd0.9	NetBSD 0.9	0.9	http://netbsd.org/netbsd/0.9
netbsd1.0	NetBSD 1.0	1.0	http://netbsd.org/netbsd/1.0
netbsd1.1	NetBSD 1.1	1.1	http://netbsd.org/netbsd/1.1
netbsd1.2	NetBSD 1.2	1.2	http://netbsd.org/netbsd/1.2
netbsd1.3	NetBSD 1.3	1.3	http://netbsd.org/netbsd/1.3
netbsd1.4	NetBSD 1.4	1.4	http://netbsd.org/netbsd/1.4
netbsd1.5	NetBSD 1.5	1.5	http://netbsd.org/netbsd/1.5
netbsd1.6	NetBSD 1.6	1.6	http://netbsd.org/netbsd/1.6
netbsd2.0	NetBSD 2.0	2.0	http://netbsd.org/netbsd/2.0
netbsd3.0	NetBSD 3.0	3.0	http://netbsd.org/netbsd/3.0
netbsd4.0	NetBSD 4.0	4.0	http://netbsd.org/netbsd/4.0
netbsd5.0	NetBSD 5.0	5.0	http://netbsd.org/netbsd/5.0
netbsd5.1	NetBSD 5.1	5.1	http://netbsd.org/netbsd/5.1
netware4	Novell Netware 4	4	http://novell.com/netware/4
netware5	Novell Netware 5	5	http://novell.com/netware/5

netware6	Novell Netware 6	6	http://novell.com/netware/6
openbsd4.2	OpenBSD 4.2	4.2	http://openbsd.org/openbsd/4.2
openbsd4.3	OpenBSD 4.3	4.3	http://openbsd.org/openbsd/4.3
openbsd4.4	OpenBSD 4.4	4.4	http://openbsd.org/openbsd/4.4
openbsd4.5	OpenBSD 4.5	4.5	http://openbsd.org/openbsd/4.5
openbsd4.8	OpenBSD 4.8	4.8	http://openbsd.org/openbsd/4.8
openbsd4.9	OpenBSD 4.9	4.9	http://openbsd.org/openbsd/4.9
openbsd5.0	OpenBSD 5.0	5.0	http://openbsd.org/openbsd/5.0
openbsd5.1	OpenBSD 5.1	5.1	http://openbsd.org/openbsd/5.1
openbsd5.2	OpenBSD 5.2	5.2	http://openbsd.org/openbsd/5.2
openbsd5.3	OpenBSD 5.3	5.3	http://openbsd.org/openbsd/5.3
openbsd5.4	OpenBSD 5.4	5.4	http://openbsd.org/openbsd/5.4
openbsd5.5	OpenBSD 5.5	5.5	http://openbsd.org/openbsd/5.5
openbsd5.6	OpenBSD 5.6	5.6	http://openbsd.org/openbsd/5.6
openbsd5.7	OpenBSD 5.7	5.7	http://openbsd.org/openbsd/5.7
openbsd5.8	OpenBSD 5.8	5.8	http://openbsd.org/openbsd/5.8
opensolaris2009.06	OpenSolaris 2009.06	2009.06	http://sun.com/opensolaris/2009.06
opensuse-factory	openSUSE	factory	http://opensuse.org/opensuse/factory
opensuse-unknown	openSUSE	unknown	http://opensuse.org/opensuse/unknown
opensuse10.2	openSUSE 10.2	10.2	http://opensuse.org/opensuse/10.2
opensuse10.3	openSUSE 10.3	10.3	http://opensuse.org/opensuse/10.3
opensuse11.0	openSUSE 11.0	11.0	http://opensuse.org/opensuse/11.0
opensuse11.1	openSUSE 11.1	11.1	http://opensuse.org/opensuse/11.1
opensuse11.2	openSUSE 11.2	11.2	http://opensuse.org/opensuse/11.2
opensuse11.3	openSUSE 11.3	11.3	http://opensuse.org/opensuse/11.3
opensuse11.4	openSUSE 11.4	11.4	http://opensuse.org/opensuse/11.4
opensuse12.1	openSUSE 12.1	12.1	http://opensuse.org/opensuse/12.1
opensuse12.2	openSUSE 12.2	12.2	http://opensuse.org/opensuse/12.2
opensuse12.3	openSUSE 12.3	12.3	http://opensuse.org/opensuse/12.3
opensuse13.1	openSUSE 13.1	13.1	http://opensuse.org/opensuse/13.1
opensuse13.2	openSUSE 13.2	13.2	http://opensuse.org/opensuse/13.2
opensuse42.1	openSUSE Leap 42.1	42.1	http://opensuse.org/opensuse/42.1
opensuse42.2	openSUSE Leap 42.2	42.2	http://opensuse.org/opensuse/42.2
opensusetumbleweed	openSUSE Tumbleweed	tumbleweed	http://opensuse.org/opensuse/tumbleweed
rhel-atomic-7.0	Red Hat Enterprise Linux Atomic Host 7.0	7.0	http://redhat.com/rhel-atomic/7.0
rhel-atomic-7.1	Red Hat Enterprise Linux Atomic Host 7.1	7.1	http://redhat.com/rhel-atomic/7.1
rhel-atomic-7.2	Red Hat Enterprise Linux Atomic Host 7.2	7.2	http://redhat.com/rhel-atomic/7.2
rhel2.1	Red Hat Enterprise Linux 2.1	2.1	http://redhat.com/rhel/2.1
rhel2.1.1	Red Hat Enterprise Linux 2.1 Update 1	2.1.1	http://redhat.com/rhel/2.1.1
rhel2.1.2	Red Hat Enterprise Linux 2.1 Update 2	2.1.2	http://redhat.com/rhel/2.1.2
rhel2.1.3	Red Hat Enterprise Linux 2.1 Update 3	2.1.3	http://redhat.com/rhel/2.1.3
rhel2.1.4	Red Hat Enterprise Linux 2.1 Update 4	2.1.4	http://redhat.com/rhel/2.1.4
rhel2.1.5	Red Hat Enterprise Linux 2.1 Update 5	2.1.5	http://redhat.com/rhel/2.1.5
rhel2.1.6	Red Hat Enterprise Linux 2.1 Update 6	2.1.6	http://redhat.com/rhel/2.1.6
rhel2.1.7	Red Hat Enterprise Linux 2.1 Update 7	2.1.7	http://redhat.com/rhel/2.1.7
rhel3	Red Hat Enterprise Linux 3	3	http://redhat.com/rhel/3
rhel3.1	Red Hat Enterprise Linux 3 Update 1	3.1	http://redhat.com/rhel/3.1
rhel3.2	Red Hat Enterprise Linux 3 Update 2	3.2	http://redhat.com/rhel/3.2
rhel3.3	Red Hat Enterprise Linux 3 Update 3	3.3	http://redhat.com/rhel/3.3
rhel3.4	Red Hat Enterprise Linux 3 Update 4	3.4	http://redhat.com/rhel/3.4
rhel3.5	Red Hat Enterprise Linux 3 Update 5	3.5	http://redhat.com/rhel/3.5
rhel3.6	Red Hat Enterprise Linux 3 Update 6	3.6	http://redhat.com/rhel/3.6
rhel3.7	Red Hat Enterprise Linux 3 Update 7	3.7	http://redhat.com/rhel/3.7
rhel3.8	Red Hat Enterprise Linux 3 Update 8	3.8	http://redhat.com/rhel/3.8
rhel3.9	Red Hat Enterprise Linux 3 Update 9	3.9	http://redhat.com/rhel/3.9
rhel4.0	Red Hat Enterprise Linux 4.0	4.0	http://redhat.com/rhel/4.0
rhel4.1	Red Hat Enterprise Linux 4.1	4.1	http://redhat.com/rhel/4.1
rhel4.2	Red Hat Enterprise Linux 4.2	4.2	http://redhat.com/rhel/4.2
rhel4.3	Red Hat Enterprise Linux 4.3	4.3	http://redhat.com/rhel/4.3
rhel4.4	Red Hat Enterprise Linux 4.4	4.4	http://redhat.com/rhel/4.4
rhel4.5	Red Hat Enterprise Linux 4.5	4.5	http://redhat.com/rhel/4.5
rhel4.6	Red Hat Enterprise Linux 4.6	4.6	http://redhat.com/rhel/4.6
rhel4.7	Red Hat Enterprise Linux 4.7	4.7	http://redhat.com/rhel/4.7
rhel4.8	Red Hat Enterprise Linux 4.8	4.8	http://redhat.com/rhel/4.8
rhel4.9	Red Hat Enterprise Linux 4.9	4.9	http://redhat.com/rhel/4.9
rhel5.0	Red Hat Enterprise Linux 5.0	5.0	http://redhat.com/rhel/5.0
rhel5.1	Red Hat Enterprise Linux 5.1	5.1	http://redhat.com/rhel/5.1
rhel5.10	Red Hat Enterprise Linux 5.10	5.10	http://redhat.com/rhel/5.10
rhel5.11	Red Hat Enterprise Linux 5.11	5.11	http://redhat.com/rhel/5.11
rhel5.2	Red Hat Enterprise Linux 5.2	5.2	http://redhat.com/rhel/5.2
rhel5.3	Red Hat Enterprise Linux 5.3	5.3	http://redhat.com/rhel/5.3
rhel5.4	Red Hat Enterprise Linux 5.4	5.4	http://redhat.com/rhel/5.4
rhel5.5	Red Hat Enterprise Linux 5.5	5.5	http://redhat.com/rhel/5.5
rhel5.6	Red Hat Enterprise Linux 5.6	5.6	http://redhat.com/rhel/5.6
rhel5.7	Red Hat Enterprise Linux 5.7	5.7	http://redhat.com/rhel/5.7
rhel5.8	Red Hat Enterprise Linux 5.8	5.8	http://redhat.com/rhel/5.8
rhel5.9	Red Hat Enterprise Linux 5.9	5.9	http://redhat.com/rhel/5.9
rhel6.0	Red Hat Enterprise Linux 6.0	6.0	http://redhat.com/rhel/6.0
rhel6.1	Red Hat Enterprise Linux 6.1	6.1	http://redhat.com/rhel/6.1
rhel6.2	Red Hat Enterprise Linux 6.2	6.2	http://redhat.com/rhel/6.2
rhel6.3	Red Hat Enterprise Linux 6.3	6.3	http://redhat.com/rhel/6.3
rhel6.4	Red Hat Enterprise Linux 6.4	6.4	http://redhat.com/rhel/6.4
rhel6.5	Red Hat Enterprise Linux 6.5	6.5	http://redhat.com/rhel/6.5
rhel6.6	Red Hat Enterprise Linux 6.6	6.6	http://redhat.com/rhel/6.6
rhel6.7	Red Hat Enterprise Linux 6.7	6.7	http://redhat.com/rhel/6.7
rhel6.8	Red Hat Enterprise Linux 6.8	6.8	http://redhat.com/rhel/6.8
rhel7.0	Red Hat Enterprise Linux 7.0	7.0	http://redhat.com/rhel/7.0
rhel7.1	Red Hat Enterprise Linux 7.1	7.1	http://redhat.com/rhel/7.1
rhel7.2	Red Hat Enterprise Linux 7.2	7.2	http://redhat.com/rhel/7.2
rhl11.0	Red Hat Linux 1.0	1.0	http://redhat.com/rhl/1.0
rhl11.1	Red Hat Linux 1.1	1.1	http://redhat.com/rhl/1.1
rhl2.0	Red Hat Linux 2.0	2.0	http://redhat.com/rhl/2.0
rhl2.1	Red Hat Linux 2.1	2.1	http://redhat.com/rhl/2.1
rhl3.0.3	Red Hat Linux 3.0.3	3.0.3	http://redhat.com/rhl/3.0.3
rhl4.0	Red Hat Linux 4.0	4.0	http://redhat.com/rhl/4.0
rhl4.1	Red Hat Linux 4.1	4.1	http://redhat.com/rhl/4.1
rhl4.2	Red Hat Linux 4.2	4.2	http://redhat.com/rhl/4.2
rhl5.0	Red Hat Linux 5.0	5.0	http://redhat.com/rhl/5.0
rhl5.1	Red Hat Linux 5.1	5.1	http://redhat.com/rhl/5.1
rhl5.2	Red Hat Linux 5.2	5.2	http://redhat.com/rhl/5.2
rhl6.0	Red Hat Linux 6.0	6.0	http://redhat.com/rhl/6.0
rhl6.1	Red Hat Linux 6.1	6.1	http://redhat.com/rhl/6.1
rhl6.2	Red Hat Linux 6.2	6.2	http://redhat.com/rhl/6.2
rhl7	Red Hat Linux 7	7	http://redhat.com/rhl/7
rhl7.1	Red Hat Linux 7.1	7.1	http://redhat.com/rhl/7.1
rhl7.2	Red Hat Linux 7.2	7.2	http://redhat.com/rhl/7.2
rhl7.3	Red Hat Linux 7.3	7.3	http://redhat.com/rhl/7.3
rhl8.0	Red Hat Linux 8.0	8.0	http://redhat.com/rhl/8.0
rhl9	Red Hat Linux 9	9	http://redhat.com/rhl/9
sled10	SUSE Linux Enterprise Desktop 10	10	http://suse.com/sled/10
sled10sp1	SUSE Linux Enterprise Desktop 10 SP1	10.1	http://suse.com/sled/10.1
sled10sp2	SUSE Linux Enterprise Desktop 10 SP2	10.2	http://suse.com/sled/10.2
sled10sp3	SUSE Linux Enterprise Desktop 10 SP3	10.3	http://suse.com/sled/10.3
sled10sp4	SUSE Linux Enterprise Desktop 10 SP4	10.4	http://suse.com/sled/10.4
sled11	SUSE Linux Enterprise Desktop 11	11	http://suse.com/sled/11

sled11sp1	SUSE Linux Enterprise Desktop 11 SP1	11.1	http://suse.com/sled/11.1
sled11sp2	SUSE Linux Enterprise Desktop 11 SP2	11.2	http://suse.com/sled/11.2
sled11sp3	SUSE Linux Enterprise Desktop 11 SP3	11.3	http://suse.com/sled/11.3
sled11sp4	SUSE Linux Enterprise Desktop 11 SP4	11.4	http://suse.com/sled/11.4
sled12	SUSE Linux Enterprise Desktop 12	12	http://suse.com/sled/12
sled12sp1	SUSE Linux Enterprise Desktop 12 SP1	12.1	http://suse.com/sled/12.1
sled12sp2	SUSE Linux Enterprise Desktop 12 SP2	12.2	http://suse.com/sled/12.2
sled9	SUSE Linux Enterprise Desktop 9	9	http://suse.com/sled/9
sles10	SUSE Linux Enterprise Server 10	10	http://suse.com/sles/10
sles10sp1	SUSE Linux Enterprise Server 10 SP1	10.1	http://suse.com/sles/10.1
sles10sp2	SUSE Linux Enterprise Server 10 SP2	10.2	http://suse.com/sles/10.2
sles10sp3	SUSE Linux Enterprise Server 10 SP3	10.3	http://suse.com/sles/10.3
sles10sp4	SUSE Linux Enterprise Server 10 SP4	10.4	http://suse.com/sles/10.4
sles11	SUSE Linux Enterprise Server 11	11	http://suse.com/sles/11
sles11sp1	SUSE Linux Enterprise Server 11 SP1	11.1	http://suse.com/sles/11.1
sles11sp2	SUSE Linux Enterprise Server 11 SP2	11.2	http://suse.com/sles/11.2
sles11sp3	SUSE Linux Enterprise Server 11 SP3	11.3	http://suse.com/sles/11.3
sles11sp4	SUSE Linux Enterprise Server 11 SP4	11.4	http://suse.com/sles/11.4
sles12	SUSE Linux Enterprise Server 12	12	http://suse.com/sles/12
sles12sp1	SUSE Linux Enterprise Server 12 SP1	12.1	http://suse.com/sles/12.1
sles12sp2	SUSE Linux Enterprise Server 12 SP2	12.2	http://suse.com/sles/12.2
sles9	SUSE Linux Enterprise Server 9	9	http://suse.com/sles/9
solaris10	Solaris 10	10	http://sun.com/solaris/10
solaris11	Oracle Solaris 11	11	http://oracle.com/solaris/11
solaris9	Solaris 9	9	http://sun.com/solaris/9
ubuntu10.04	Ubuntu 10.04 LTS	10.04	http://ubuntu.com/ubuntu/10.04
ubuntu10.10	Ubuntu 10.10	10.10	http://ubuntu.com/ubuntu/10.10
ubuntu11.04	Ubuntu 11.04	11.04	http://ubuntu.com/ubuntu/11.04
ubuntu11.10	Ubuntu 11.10	11.10	http://ubuntu.com/ubuntu/11.10
ubuntu12.04	Ubuntu 12.04 LTS	12.04	http://ubuntu.com/ubuntu/12.04
ubuntu12.10	Ubuntu 12.10	12.10	http://ubuntu.com/ubuntu/12.10
ubuntu13.04	Ubuntu 13.04	13.04	http://ubuntu.com/ubuntu/13.04
ubuntu13.10	Ubuntu 13.10	13.10	http://ubuntu.com/ubuntu/13.10
ubuntu14.04	Ubuntu 14.04 LTS	14.04	http://ubuntu.com/ubuntu/14.04
ubuntu14.10	Ubuntu 14.10	14.10	http://ubuntu.com/ubuntu/14.10
ubuntu15.04	Ubuntu 15.04	15.04	http://ubuntu.com/ubuntu/15.04
ubuntu15.10	Ubuntu 15.10	15.10	http://ubuntu.com/ubuntu/15.10
ubuntu16.04	Ubuntu 16.04	16.04	http://ubuntu.com/ubuntu/16.04
ubuntu4.10	Ubuntu 4.10	4.10	http://ubuntu.com/ubuntu/4.10
ubuntu5.04	Ubuntu 5.04	5.04	http://ubuntu.com/ubuntu/5.04
ubuntu5.10	Ubuntu 5.10	5.10	http://ubuntu.com/ubuntu/5.10
ubuntu6.06	Ubuntu 6.06 LTS	6.06	http://ubuntu.com/ubuntu/6.06
ubuntu6.10	Ubuntu 6.10	6.10	http://ubuntu.com/ubuntu/6.10
ubuntu7.04	Ubuntu 7.04	7.04	http://ubuntu.com/ubuntu/7.04
ubuntu7.10	Ubuntu 7.10	7.10	http://ubuntu.com/ubuntu/7.10
ubuntu8.04	Ubuntu 8.04 LTS	8.04	http://ubuntu.com/ubuntu/8.04
ubuntu8.10	Ubuntu 8.10	8.10	http://ubuntu.com/ubuntu/8.10
ubuntu9.04	Ubuntu 9.04	9.04	http://ubuntu.com/ubuntu/9.04
ubuntu9.10	Ubuntu 9.10	9.10	http://ubuntu.com/ubuntu/9.10
win1.0	Microsoft Windows 1.0	1.0	http://microsoft.com/win/1.0
win10	Microsoft Windows 10	10.0	http://microsoft.com/win/10
win2.0	Microsoft Windows 2.0	2.0	http://microsoft.com/win/2.0
win2.1	Microsoft Windows 2.1	2.1	http://microsoft.com/win/2.1
win2k	Microsoft Windows 2000	5.0	http://microsoft.com/win/2k
win2k12	Microsoft Windows Server 2012	6.3	http://microsoft.com/win/2k12
win2k12r2	Microsoft Windows Server 2012 R2	6.3	http://microsoft.com/win/2k12r2
win2k3	Microsoft Windows Server 2003	5.2	http://microsoft.com/win/2k3
win2k3r2	Microsoft Windows Server 2003 R2	5.2	http://microsoft.com/win/2k3r2
win2k8	Microsoft Windows Server 2008	6.0	http://microsoft.com/win/2k8
win2k8r2	Microsoft Windows Server 2008 R2	6.1	http://microsoft.com/win/2k8r2
win3.1	Microsoft Windows 3.1	3.1	http://microsoft.com/win/3.1
win7	Microsoft Windows 7	6.1	http://microsoft.com/win/7
win8	Microsoft Windows 8	6.2	http://microsoft.com/win/8
win8.1	Microsoft Windows 8.1	6.3	http://microsoft.com/win/8.1
win95	Microsoft Windows 95	4.0	http://microsoft.com/win/95
win98	Microsoft Windows 98	4.1	http://microsoft.com/win/98
winme	Microsoft Windows Millennium Edition	4.9	http://microsoft.com/win/me
winnt3.1	Microsoft Windows NT Server 3.1	3.1	http://microsoft.com/winnt/3.1
winnt3.5	Microsoft Windows NT Server 3.5	3.5	http://microsoft.com/winnt/3.5
winnt3.51	Microsoft Windows NT Server 3.51	3.51	http://microsoft.com/winnt/3.51
winnt4.0	Microsoft Windows NT Server 4.0	4.0	http://microsoft.com/winnt/4.0
winvista	Microsoft Windows Vista	6.0	http://microsoft.com/win/vista
winxp	Microsoft Windows XP	5.1	http://microsoft.com/win/xp

Listing the Debian GNU/Linux options.

ada:~\$ osinfo-query os grep debian			
debian1.1	Debian Buzz	1.1	http://debian.org/debian/1.1
debian1.2	Debian Rex	1.2	http://debian.org/debian/1.2
debian1.3	Debian Bo	1.3	http://debian.org/debian/1.3
debian2.0	Debian Hamm	2.0	http://debian.org/debian/2.0
debian2.1	Debian Slink	2.1	http://debian.org/debian/2.1
debian2.2	Debian Potato	2.2	http://debian.org/debian/2.2
debian3	Debian Woody	3	http://debian.org/debian/3
debian3.1	Debian Sarge	3.1	http://debian.org/debian/3.1
debian4	Debian Etch	4	http://debian.org/debian/4
debian5	Debian Lenny	5	http://debian.org/debian/5
debian6	Debian Squeeze	6	http://debian.org/debian/6
debian7	Debian Wheezy	7	http://debian.org/debian/7
debian8	Debian Jessie	8	http://debian.org/debian/8
debian9	Debian Stretch	9	http://debian.org/debian/9
debiantesting	Debian Testing	testing	http://debian.org/debian/testing

6.4 Build first domain

As an example build a Debian 8.7 64-bit domain. This is achieved using *virt-install* which is a tool for creating new KVM, Xen, or GNU/Linux container guests using the "libvirt" hypervisor management library.

```
ada:~$ virt-install \
    --connect qemu:///system \
    --virt-type=kvm \
    --name vm01 \
    --ram 512 \
    --vcpus=2 \
    --disk path=/virt/kvm/images/vm01.img,size=12 \
    --cdrom /virt/iso/Debian-8.7.1/debian-8.7.1-amd64.iso \
    --graphics vnc,listen=0.0.0.0 \
    --noautoconsole \
    --os-type linux \
    --os-variant debian8 \
    --network=bridge:br0 \
    --hvm

Starting install...
Allocating 'vm01.img' | 12 GB 00:00:00
Creating domain... | 0 B 00:00:00
Domain installation still in progress. You can reconnect to
the console to complete the installation process.
```

- **connect** - Connect to the hypervisor.
 - **qemu:///system** - For creating KVM/QEMU guests to be run by the system libvirtd instance.
- **virt-type** - The hypervisor to install on. Example choices are kvm, qemu, xen, or kqemu.
- **name** - Name of the new guest virtual machine instance.
- **ram** - Random Access Memory (RAM) allocated for the guest, in megabytes.
- **disk** - Specifies media to use as storage for the guest.
 - **size** - Size (in Gigabytes (GB)) to use if creating new storage.
- **cdrom** - Source of installation, an ISO file.
- **graphics** - Setup either a VNC or SPICE Server in the host that allows access.
 - **listen** - Defines either 127.0.0.1 (local access), 0.0.0.0 (global access or a specific IP address for access).
- **noautoconsole** - Don't automatically try to connect to the guest console.
- **os-type** - Optimise the guest configuration for a type of operating system (ex. *linux*, *windows*)
- **os-variant** - Further optimise the guest configuration for a specific OS (if this is used *os-type* is not strictly necessary).
- **network** - Connect the guest to the host network. (use **brctl show** to see)
- **hvm** - Request the use of full virtualisation, if both ParaVirtualisation (PV) & full virtualisation are available on the host.

Note there is an option to use either the SPICE or VNC remote desktop protocol. The *virt-viewer* will connect to whichever is configured in the domain at build time. There is also an option to change a domain with the *virt-manager* domain details which will be seen later. Essentially the SPICE protocol is generally faster than VNC.

6.5 Connect to a domain

To connect to the domain use the *virt-viewer* tool and either identify the VM by its *Id*, in this case *1* or by its *name* which in this case is *vm01*.

```
ada:~$ virt-viewer --connect qemu:///system 1
```

or

```
ada:~$ virt-viewer --connect qemu:///system vm01
```

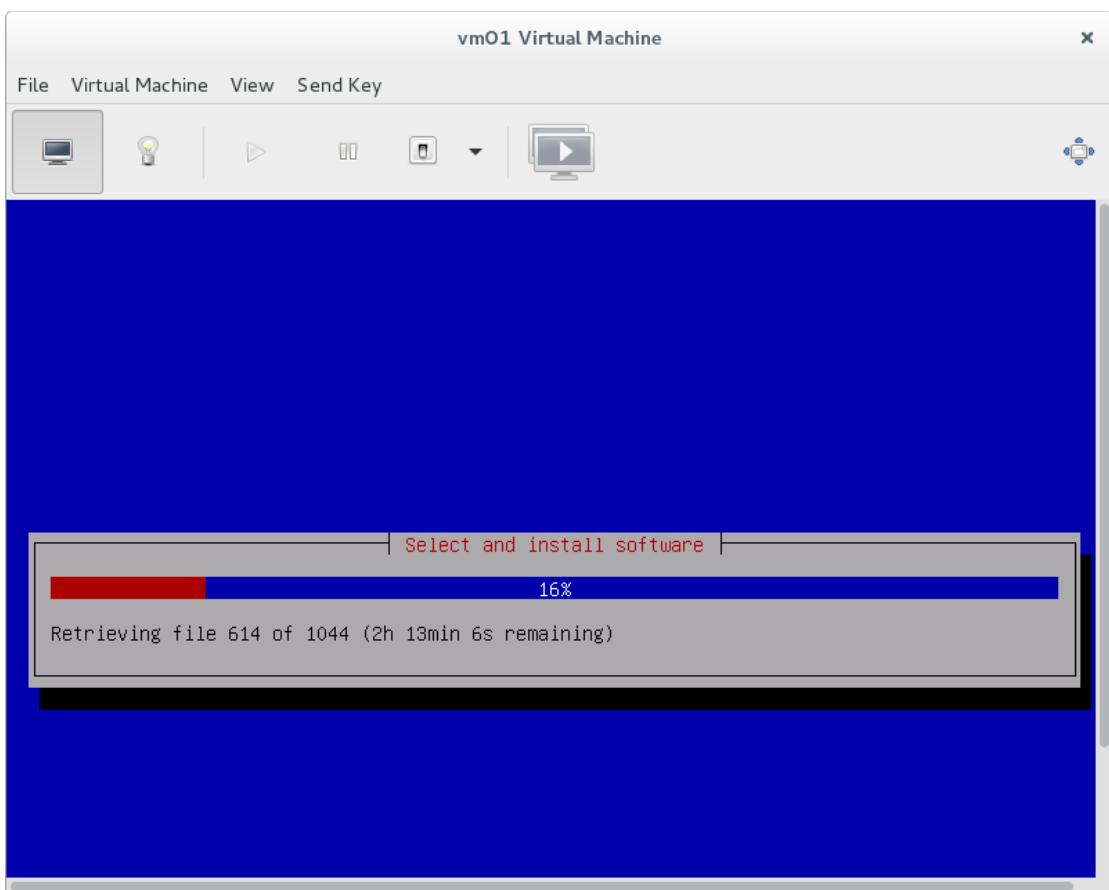


Illustration 4: Debian install

Follow the Debian install script and the domain guest will install.

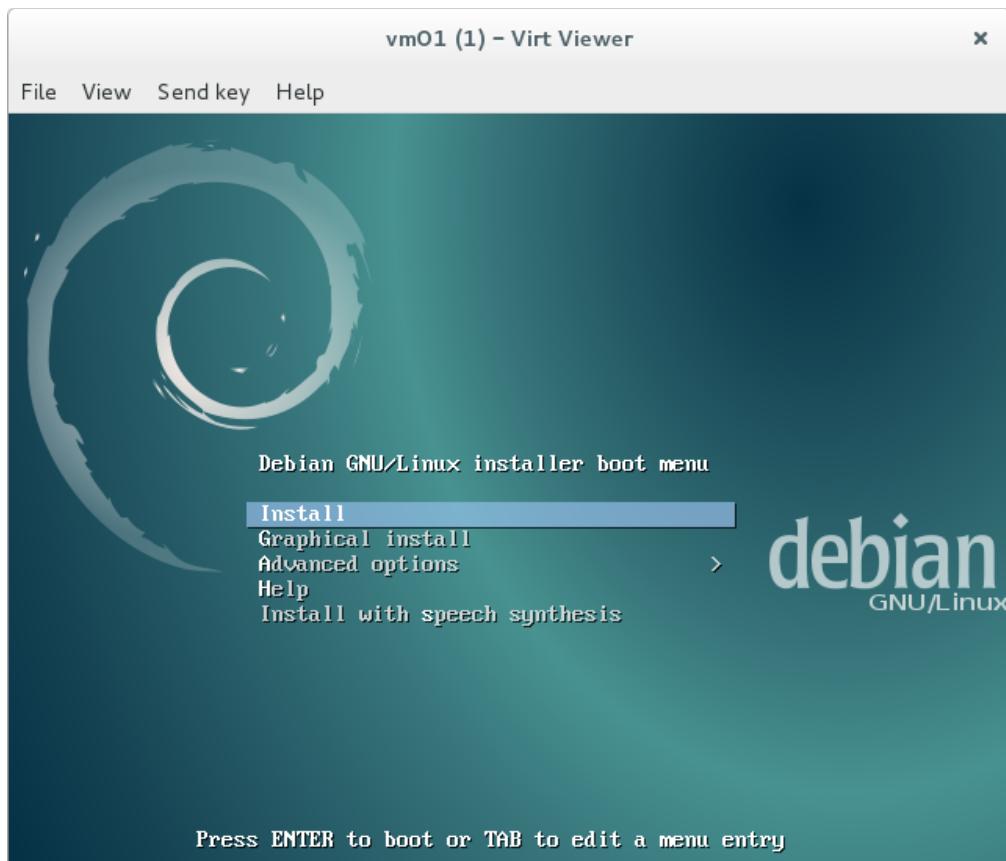


Illustration 5: Debian guest

6.6 Build a Microsoft Windows domain

As a second example build a Microsoft Windows 7, 64-bit domain.

```
ada:~$ virt-install \
--connect qemu:///system \
--virt-type=kvm \
--name vm02 \
--ram 1024 \
--vcpus=2 \
--disk path=/virt/kvm/images/vm02.img,size=12 \
--cdrom /virt/iso/Windows-7/Microsoft_Windows_7.iso \
--graphics spice,listen=0.0.0.0 \
--noautoconsole \
--os-type windows \
--os-variant win7 \
--network=bridge:br0 \
--hvm
```

```
Starting install...
Allocating 'vm02.img' | 12 GB    00:00
Creating domain...     | 0 B      00:00
Domain installation still in progress. Waiting for
installation to complete.
```

Connect to the running installation graphic.

```
ada:~$ virt-viewer --connect qemu:///system vm02
```



Illustration 6: Windows guest install

Follow the Microsoft Windows install process, it will reboot a number of times as part of the install which requires the *virt-viewer* to be restarted to reconnect.

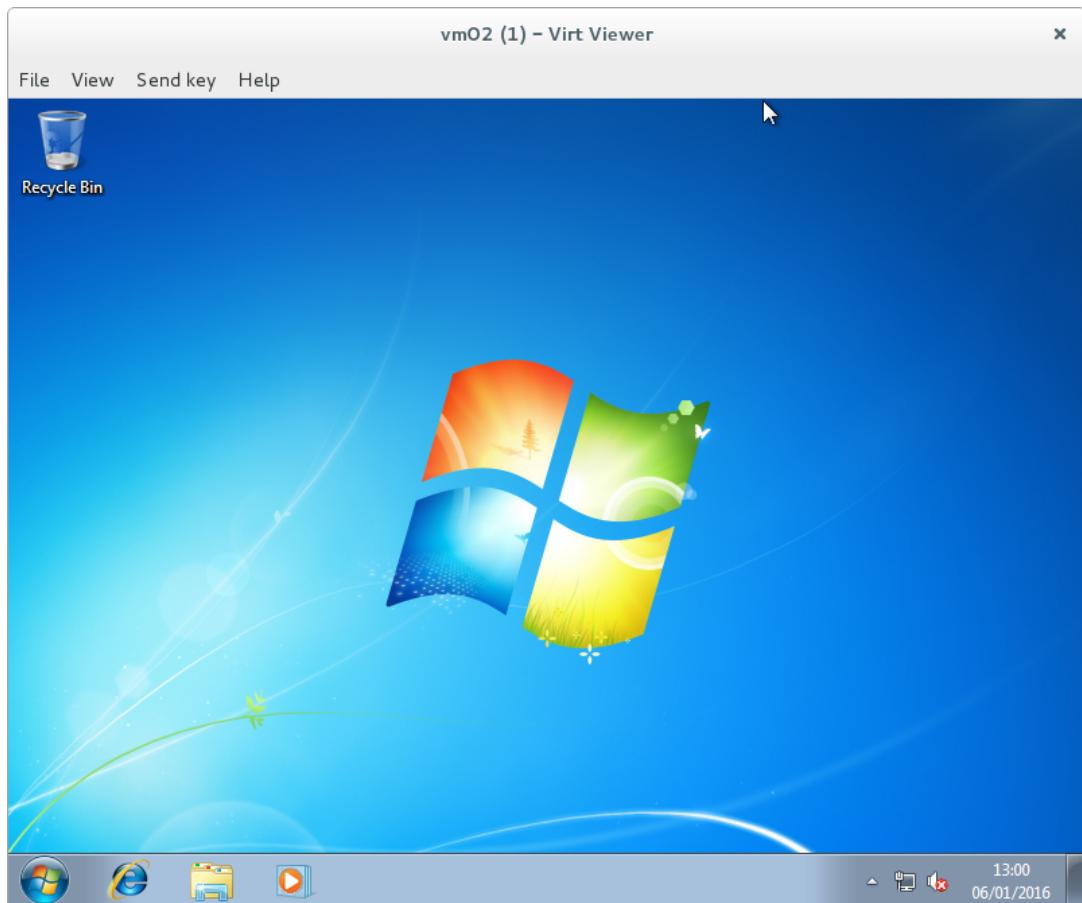


Illustration 7: Windows guest

6.7 Build a domain over the Internet

As an example build a Ubuntu 64-bit domain. This is achieved using *virt-install* which is a tool for creating new KVM, Xen, or GNU/Linux container guests using the *libvirt* hypervisor management library.

```
ada:~$ virt-install \
--connect qemu:///system \
--virt-type=kvm \
--name vm03 \
--ram 512 \
--vcpus=2 \
--disk path=/virt/kvm/images/vm03.img,size=12 \
--location
'http://archive.ubuntu.com/ubuntu/dists/wily/main/installer-
amd64/' \
--graphics spice,listen=0.0.0.0 \
--noautoconsole \
--os-type linux \
--os-variant ubuntusaucy \
--network=bridge:br0 \
--hvm
```

```
Starting install...
Retrieving file MANIFEST...
Retrieving file linux...          21% [=====
Retrieving file linux...          21% [=====
Retrieving file linux...          | 13 MB   00:55 ...
Retrieving file initrd.gz...     | 44 MB   03:12 ...
Allocating 'virtinst-linux.vj1SiL'| 6.5 MB   00:00
Transferring virtinst-linux.vj1SiL| 6.5 MB   00:00
Allocating 'virtinst-initrd.gz.FkqiPN'| 22 MB   00:00
Transferring virtinst-initrd.gz.FkqiPN| 22 MB   00:00
Allocating 'vm03.img'            | 12 GB   00:00
Creating domain...               | 0 B    00:00
Domain installation still in progress. You can reconnect to
the console to complete the installation process.
```

- **location** - Used instead of *cdrom* when point to an image on the Internet. Note that the link specified is to the root directory and not to the *.iso* image itself.

7. Connecting to Virtual Machines

7.1 Connect to a remote domain

Install a SPICE or VNC client.

7.1.1 SPICE

```
ada:~$ sudo apt-get install spice-client
ada:~$ sudo apt-get install spice-client-gtk
ada:~$ sudo apt-get install python-spice-client-gtk
```

Find the domain port number on the hypervisor and confirm it is configured for SPICE.

```
virsh # domdisplay vm05
spice://localhost:5905
```

Connect with the client.

```
ada:~$ spicec --host 213.74.34.100 --port 5905
```

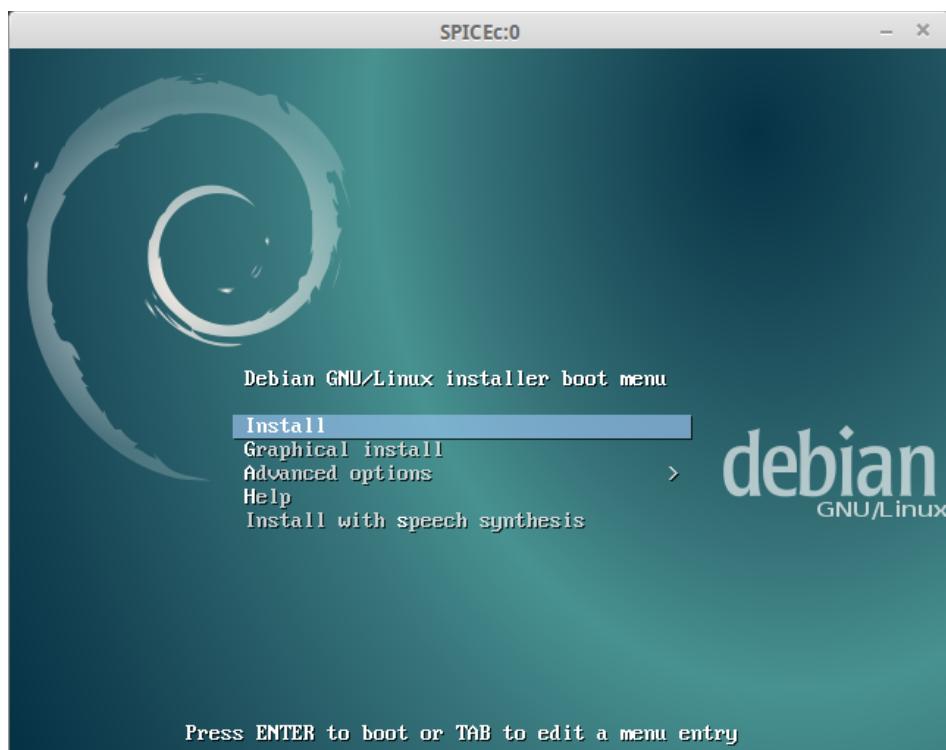


Illustration 8: Debian guest through SPICE viewer

For Microsoft Windows domains it will be necessary to install Windows guest tools. This can be downloaded from <http://www.spice-space.org/download.html>.

7.1.2 VNC

Install a VNC client.

```
ada:~$ sudo apt-get install vncviewer
```

Find the domain port number on the hypervisor and confirm it is configured for VNC.

```
virsh # domdisplay vm02
vnc://localhost:8
```

Connect with the client

```
ada:~$ vncviewer 213.74.34.100:8
```

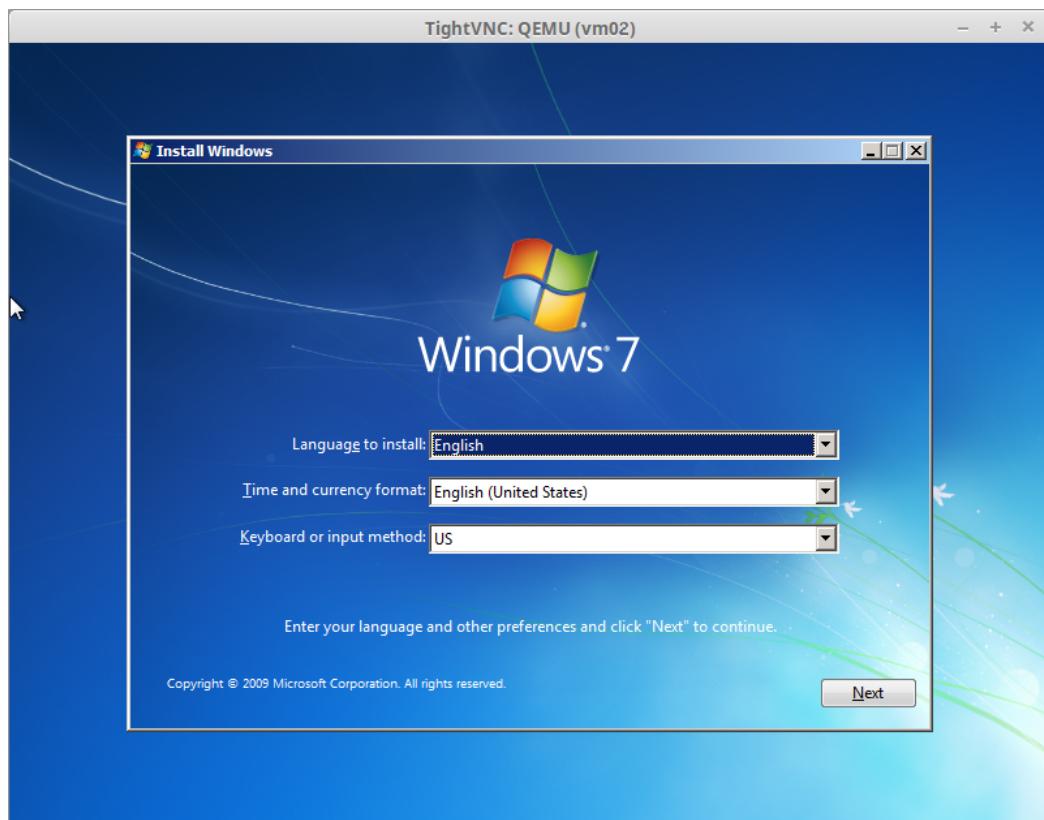


Illustration 9: Windows guest through VNC viewer

8. Manage and operate the domains

8.1 Check domain list

```
ada:~$ virsh --connect qemu:///system

virsh # list
  Id   Name          State
  --  --
  1   vm01         running
  2   vm02         running
  3   vm03         running

virsh # list --all
  Id   Name          State
  --  --
  1   vm01         running
  -   vm02        shut off
  3   vm03         running
```

Some notes on possible states.

- **running** - domains which are currently active on a CPU.
- **blocked / blocking** - domain is presently idle, waiting for I/O or waiting for the hypervisor.
- **paused** - domain is suspended.
- **shutdown** - domains in the process of shutting down.
- **Domains** - domain is off and not using system resources.
- **crashed** - domain failed while running and is no longer running.

Displaying domain information

```
virsh # dominfo vm01
Id:                 2
Name:               vm01
UUID:              b4a8695a-ccb-49b1-9acd-cd4f649285ba
OS Type:            hvm
State:              running
CPU(s):             2
CPU time:           244.3s
Max memory:         524288 KiB
Used memory:        524288 KiB
Persistent:         yes
Autostart:          disable
Managed save:       no
Security model:     none
Security DOI:      0
```

8.2 Shutdown a domain

```
virsh # shutdown vm01
Domain vm01 is being shutdown
```

8.3 Reboot a domain

```
virsh # reboot vm01
Domain vm01 is being rebooted
```

8.4 Terminate a domain

Destroying a domain is an immediate ungraceful shutdown and stops any guest domain sessions. Use the *destroy* option only when the guest is non-responsive.

```
ada:~$ virsh --connect qemu:///system

virsh # destroy vm01
Domain vm01 destroyed
```

8.5 Start a domain

```
virsh # start vm01
Domain vm01 started
```

8.6 Suspend a domain

A running domain maybe suspended. The domain is kept in memory but it will not be scheduled any more and therefore it still consumes system RAM. Disk and network I/O will not occur while the guest is suspended. This operation is immediate and the guest must be restarted with the *resume* option.

```
virsh # suspend vm02
Domain vm02 suspended

virsh # list --all
  Id   Name           State
  --  --
  1    vm01          running
  2    vm02          paused
  3    vm03          running
```

8.6.1 Resume a domain

Resume a suspended domain.

```
virsh # resume vm02
Domain vm02 resumed

virsh # list --all
  Id      Name           State
  --
  1      vm01          running
  2      vm02          running
  3      vm03          running
```

8.7 Save a domain

Saving a domain stops the guest VM and saves the data to a file, which may take some time given the amount of memory in use by the domain. The state of the guest can be recovered with the *restore* option.

```
virsh # save vm02 vm02file.vmsav

Domain vm02 saved to vm02file.vmsav

virsh # list --all
  Id      Name           State
  --
  1      vm01          running
  3      vm03          running

/virt/kvm/images$ ls
vm01.img  vm02file.vmsav  vm02.img
```

8.7.1 Restoring a domain

To restore a guest that was previously saved.

```
virsh # restore /virt/kvm/images/vm02file.vmsav
Domain restored from /virt/kvm/images/vm02file.vmsav

virsh # list --all
  Id      Name           State
  --
  1      vm01          running
  2      vm02          running
  3      vm03          running
```

8.8 Cloning a domain

First the domain must be *suspended* or *shutdown* prior to cloning.

```
virsh # suspend vm01
Domain vm01 suspended

virsh # exit

ada:~$ virt-clone --connect qemu:///system \
      --original vm01 \
      --name cloned-debian-vm \
      --file /virt/kvm/images/cloned-debian-vm

WARNING Setting the graphics device port to autoport, in order
to avoid conflicting.
Allocating 'cloned-debian-vm' | 12 GB 00:49

Clone 'cloned-debian-vm' created successfully.
```

Now that it is created check it is there and run it.

```
virsh # list --all
      Id      Name           State
-----
      1      vm01            paused
      2      vm02            running
      3      vm03            running
      -      cloned-debian-vm    shut off

virsh # start cloned-debian-vm
Domain cloned-debian-vm started

virsh # list --all
      Id      Name           State
-----
      1      vm01            paused
      2      vm02            running
      3      vm03            running
      4      cloned-debian-vm    running
```

Now resume the original *vm01* and check that all four domains are in fact running.

```
virsh # resume vm01
Domain vm01 resumed

virsh # list --all
      Id      Name           State
-----
      1      vm01            running
      2      vm02            running
      3      vm03            running
      4      cloned-debian-vm    running
```

8.9 Delete a domain

Firstly *destroy*, then *undefine* it and *vol-delete* the associated volume to completely remove the domain VM.

```
virsh # destroy vm01
Domain vm01 has been destroyed

virsh # undefine vm01
Domain vm01 has been undefined

virsh # vol-delete /virt/kvm/images/vm01.img
Vol /virt/kvm/images/vm01.img deleted
```

8.10 Domain XML file

This XML file defines the detail of the domain VM. It is possible to edit this file to implement change and reload the domain to incorporate these changes. As an example change the name from *vm01* to *vm01_Debian_8* and reload.

```
ada:~$ virsh --connect qemu:///system dumpxml vm01 > vm01.xml

ada:~$ cat vm01.xml
<domain type='kvm' id='1'>
  <name>vm01</name>
  <uuid>b4a8695a-ccb-49b1-9acd-cd4f649285ba</uuid>
  <memory unit='KiB'>524288</memory>
  <currentMemory unit='KiB'>524288</currentMemory>
  <vcpu placement='static'>2</vcpu>
  <resource>
    <partition>/machine</partition>
  </resource>
  <os>
    <type arch='x86_64' machine='pc-i440fx-2.1'>hvm</type>
    <boot dev='hd' />
  </os>
  <features>
    <acpi/>
    <apic/>
    <pae/>
  </features>
  <cpu mode='custom' match='exact'>
    <model fallback='allow'>Haswell</model>
  </cpu>
  <clock offset='utc'>
    <timer name='rtc' tickpolicy='catchup' />
    <timer name='pit' tickpolicy='delay' />
    <timer name='hpet' present='no' />
  </clock>
  <on_poweroff>destroy</on_poweroff>
  <on_reboot>restart</on_reboot>
  <on_crash>restart</on_crash>
  <devices>
    <emulator>/usr/bin/kvm</emulator>
    <disk type='file' device='disk'>
      <driver name='qemu' type='qcow2' />
      <source file='/virt/kvm/images/vm01.snapshot01' />
      <backingStore type='file' index='1'>
        <format type='raw' />
        <source file='/virt/kvm/images/vm01.img' />
      </backingStore>
    </disk>
    <disk type='block' device='cdrom'>
      <driver name='qemu' type='raw' />
    </disk>
  </devices>

```

```

<target dev='hda' bus='ide' />
<readonly/>
<alias name='ide0-0-0-0' />
<address type='drive' controller='0' bus='0' target='0' unit='0' />
</disk>
<controller type='usb' index='0' model='ich9-ehci1' >
  <alias name='usb0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x7' />
</controller>
<controller type='usb' index='0' model='ich9-uhci1' >
  <alias name='usb0' />
  <master startport='0' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0' multifunction='on' />
  </controller>
  <controller type='usb' index='0' model='ich9-uhci2' >
    <alias name='usb0' />
    <master startport='2' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x1' />
  </controller>
  <controller type='usb' index='0' model='ich9-uhci3' >
    <alias name='usb0' />
    <master startport='4' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x2' />
  </controller>
  <controller type='pci' index='0' model='pci-root' >
    <alias name='pci.0' />
  </controller>
  <controller type='ide' index='0' >
    <alias name='ide0' />
    <address type='pci' domain='0x0000' bus='0x00' slot='0x01' function='0x1' />
  </controller>
<interface type='bridge' >
  <mac address='52:54:00:c5:50:05' />
  <source bridge='br0' />
  <target dev='vnet0' />
  <model type='virtio' />
  <alias name='net0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />
</interface>
<serial type='pty' >
  <source path='/dev/pts/0' />
  <target port='0' />
  <alias name='serial0' />
</serial>
<console type='pty' tty='/dev/pts/0' >
  <source path='/dev/pts/0' />
  <target type='serial' port='0' />
  <alias name='serial0' />
</console>
<input type='tablet' bus='usb' >
  <alias name='input0' />
</input>
<input type='mouse' bus='ps2' />
<input type='keyboard' bus='ps2' />
<graphics type='vnc' port='5900' autoport='yes' listen='127.0.0.1' >
  <listen type='address' address='127.0.0.1' />
</graphics>
<video>
  <model type='cirrus' vram='9216' heads='1' />
  <alias name='video0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x02' function='0x0' />
</video>
<memballoon model='virtio' >
  <alias name='balloon0' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x06' function='0x0' />
</memballoon>
</devices>
</domain>

```

Using the sed stream editor, make the change and backup the original file.

```
ada:~$ sed -i.bak 's/<name>vm01</<name>vm01_Debian_8</' vm01.xml
```

Confirm the change.

```
ada:~$ diff vm01.xml vm01.xml.bak
2c2
<    <name>vm01_Debian_8</name>
---
>    <name>vm01</name>
```

undefine the old domain to prevent an error resulting from a duplicate UUID.

```
ada:~$ virsh --connect qemu:///system

virsh # undefine vm01
Domain vm01 has been undefined

virsh # destroy vm01
virsh # start vm01_Debian_8

virsh # define vm01.xml
Domain vm01_Debian_8 defined from vm01.xml

virsh # start vm01_Debian_8
Domain vm01_Debian_8 started
```

8.11 Creating a snapshot

A snapshot is a copy of the domain disk file at a given point in time. Snapshots are used to restore a domain to a particular point in time when a failure or system error occurs.

There are two classes of snapshots for QEMU guests. Internal snapshots are contained completely within a QEMU Copy On Write version 2 (QCOW2) file but have the drawback of slow creation times, less maintenance by upstream QEMU, and the requirement to use QCOW2 disks.

External snapshots are more useful, because they work with any type of original disk image, can be taken without guest downtime. They are created when using the *virsh snapshot-create-as --disk-only* option. Unfortunately until improvements are made in *libvirt*, snapshots can only be made, they cannot be restored through *libvirt*. However there is a workaround using the eXtensible Markup Language (XML) file for the domain.

Here are the domains running

virsh # list		
Id	Name	State
1	vm01	running
2	vm02	running
3	vm03	running

Looking at the first domain *vm01*, confirm that there are no existing snapshots.

```
virsh # snapshot-list vm01
Name           Creation Time           State
-----
```

Using *virsh* create a snapshot, note that the snapshot will be in QCOW format and not RAW.

```
virsh # snapshot-create-as --domain vm01 --name "snap01-vm01"
        --description "First snapshot" --atomic
Domain snapshot snap01-vm01 created
```

The following options can be used for more granularity.

- **--diskspec** - adds the disk elements to the Snapshot XML file.
 - **--diskspec vda, file=/virt/kvm/images/snap01-vm01.img**
- **--disk-only** - takes the snapshot of only the disk.
- **--atomic** - either the snapshot is run completely or fails w/o making any changes.

Review the new snapshot.

```
virsh # snapshot-list vm01
Name           Creation Time           State
-----
snap01-vm01   2016-01-11 12:14:10 +0000 disk-snapshot
```

Review the files in the directory and the file types. Note that the snapshot is of QEMU Copy On Write version 3 (QCOW3) format.

```
ada:~$ ls /virt/kvm/images/*vm01*
/virt/kvm/images/snap01-vm01.img  /virt/kvm/images/vm01.img

ada:~$ sudo file /virt/kvm/images/vm01.img
/virt/kvm/images/vm01.img: DOS/MBR boot sector

ada:~$ sudo file /virt/kvm/images/snap01-vm01.img
/virt/kvm/images/snap01-vm01.img: QEMU QCOW Image (v3), has
backing    file (path /virt/kvm/images/vm01.img), 12582969344
bytes

ada:~$ sudo qemu-img info /virt/kvm/images/vm01.img
image: /virt/kvm/images/vm01.img
file format: raw
virtual size: 12G (12582969344 bytes)
disk size: 12G
```

```
ada:~$ sudo qemu-img info /virt/kvm/images/snap01-vm01.img
image: /virt/kvm/images/snap01-vm01.img
file format: qcow2
virtual size: 12G (12582969344 bytes)
disk size: 5.7M
cluster_size: 65536
backing file: /virt/kvm/images/vm01.img
backing file format: raw
Format specific information:
    compat: 1.1
    lazy refcounts: false
```

Confirm the image that the domain is running currently.

```
virsh # domblklist vm01
Target      Source
-----
vda        /virt/kvm/images/snap01-vm01.img
hdc        -
```

Make some changes to the running domain. Add a file for example.

```
vm01~$ echo 'I am a file added to the snapshot' >>
~/snapshot.txt
```

8.12 Revert to a previous snapshot

It is possible to revert by one of two methods. Using the *virsh* command *snapshot-revert* *vm01 snap01-vm01.img* or by editing the domain XML file.

Shutdown the domain, a simple shutdown of the guest OS is fine. And confirm it is in fact not running.

```
virsh # list --all
Id      Name           State
-----
2       vm02          running
3       vm03          running
-       vm01          shut off
```

8.12.1 snapshot-revert method

The *virsh* command *snapshot-revert* didn't work on earlier versions of KVM. Here is an example of how it currently works.

```
snapshot-revert <domain> [--snapshotname <string>] [--current]
[--running] [--paused] [--force]

--domain <string> domain name, id or uuid
--snapshotname <string> snapshot name
--current      revert to current snapshot
--running     after reverting, change state to running
--paused      after reverting, change state to paused

virsh # snapshot-revert --domain vm01 --snapshotname snap01-
vm01 --running
```

8.12.2 Edit the domain file method

Edit the domain XML file to change the name of the image back to the original.

```
virsh # edit vm01
```

Change the source file to the original image and change the type to raw to match the original image.

```
....  
....  
  
<disk type='file' device='disk'>  
  <driver name='qemu' type='qcow2'/>  
  <source file='/virt/kvm/images/snap01-vm01.img'/'>  
  <target dev='vda' bus='virtio'/'>  
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/'>  
</disk>  
  
....  
....  
to:  
  
....  
....  
<disk type='file' device='disk'>  
  <driver name='qemu' type='raw'/'>  
  <source file='/virt/kvm/images/vm01.img'/'>  
  <target dev='vda' bus='virtio'/'>  
  <address type='pci' domain='0x0000' bus='0x00' slot='0x04' function='0x0'/'>  
</disk>  
  
....  
....
```

The change is confirmed back to the console.

```
Domain vm01 XML configuration edited.
```

Now start the domain.

```
virsh # start vm01  
Domain vm01 started
```

Check the running domain block device. It has reverted to the original image.

```
virsh # domblklist vm01  
Target      Source  
-----  
vda          /virt/kvm/images/vm01.img  
hdc          -
```

Check and confirm that the *snapshot.txt* file is there.

```
vm01~$ ls snapshot.txt  
ls: cannot access snapshot.txt: No such file or directory
```

8.13 Delete a snapshot

If a snapshot is to be deleted, remove the metadata defining it and then delete the file.

```
virsh # snapshot-delete --domain vm01 --metadata "snap01-vm01"  
Domain snapshot snap01-vm01 deleted  
  
ada:~$ sudo rm /virt/kvm/image/snap01-vm01.img
```

This page is intentionally blank

9. Virtual networks

It is possible to build virtual networks within the KVM Hypervisor and even link them to the outside physical network. By way of explanation the document will go through the steps to build the following network. All elements are in fact virtual. *vm01* will act as a router and will have access to the physical world via a bridged connection to *br0* which has a connection to *eth0* on the host computer. Like has been shown already this receives an IP address from a Dynamic Host Configuration Protocol (DHCP) Server on the network. The router will have a second virtual interface *eth1* which is connected to a virtual network *pvtbr0* and has a static IP address 10.1.1.1/24 which will act as an IP gateway for the other guests.

A Windows 7 guest *vm02*, and two Debian 8 guests *vm03* plus *vm04* are configured with their primary interfaces connected to the *pvtbr0* virtual network, with IP addresses 10.1.1.2, 10.1.1.3 and 10.1.1.4 respectfully. Each has 10.1.1.1 as its gateway.

The router *vm01* is configured to forward IP traffic and has a *iptables* masquerade rule to give the guests access to the Internet in the physical world via *Network Address Translation (NAT)*.

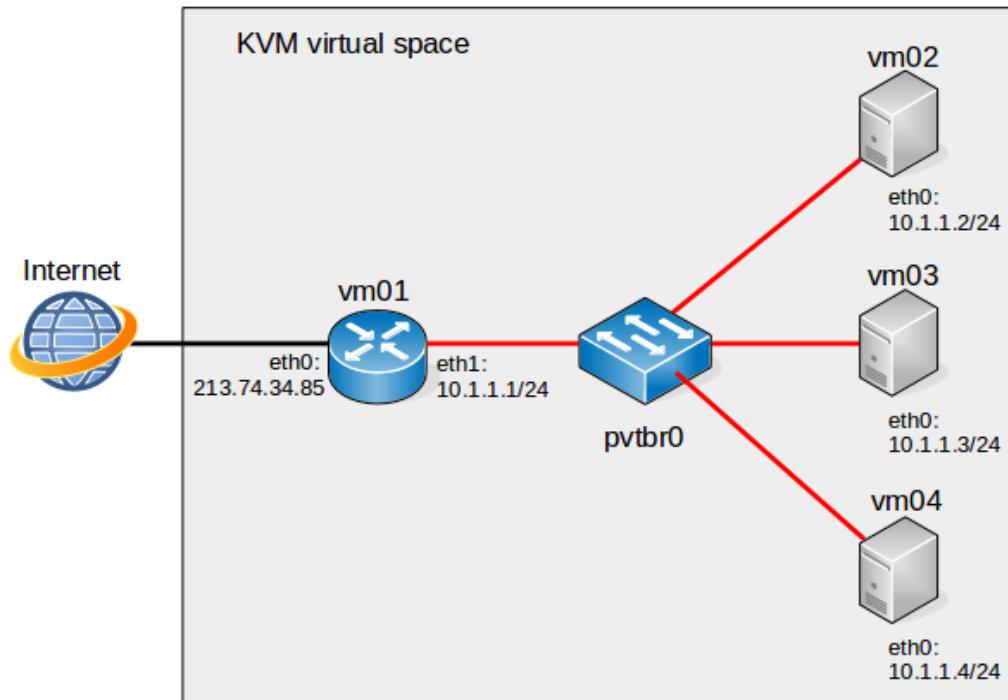


Illustration 10: Virtual network

9.1 Create a new network

With a number of guest domains on a KVM host it may become desirable to have a virtual network. List existing networks, presuming there is only one at this stage.

```
ada:~$ virsh --connect qemu:///system net-list --all
Name           State   Autostart  Persistent
-----
default        inactive  no        yes
```

Create a new network XML file.

```
ada:~$ sudo bash -c 'cat << EOM >> /virt/kvm/networks/pvtnet01.xml
<network>
  <name>pvtnet01</name>
  <bridge name="pvtbr0" />
</network>
EOM'
```

```
ada:~$ cat /virt/kvm/networks/pvtnet01.xml
<network>
  <name>pvtnet01</name>
  <bridge name="pvtbr0" />
</network>
```

Define the network in the KVM hypervisor.

```
virsh # net-define /virt/kvm/networks/pvtnet01.xml
Network pvtnet01 defined from /virt/kvm/networks/pvtnet01.xml

virsh # net-list --all
Name           State   Autostart  Persistent
-----
default        inactive  no        yes
pvtnet01      inactive  no        yes
```

start and *enable* the automatic startup of the new virtual network.

```
virsh # net-start pvtnet01
Network pvtnet01 started

virsh # net-autostart pvtnet01
Network pvtnet01 marked as autostarted

virsh # net-list --all
Name           State   Autostart  Persistent
-----
default        inactive  no        yes
pvtnet01      active   yes       yes
```

The bridge created can be seen on the hypervisor host OS.

```
ada:~$ sudo brctl show
bridge name      bridge id      STP enabled      interfaces
br0              8000.305a3a083921    no               eth0
                                         vnet2
pvtbr0           8000.5254002db430    yes              pvtbr0-nic
                                         vnet0
```

9.2 Configure domain VMs to connect to the new network

Edit each domain *vm02*, *vm03* and *vm04* by replacing the existing *<interface>* as shown. A choice of *nano* or *Vi IMproved (VIM)* editors is given the first time the *edit* command is used. Here is an example of one domain being configured.

```
virsh # shutdown vm03
Domain vm03 is being shutdown

virsh # domstate vm03
shut off

virsh # edit vm03
```

Change:

```
....
....
<interface type='bridge'>
  <mac address='52:54:00:c5:50:05' />
  <source bridge='br0' />
  <model type='virtio' />
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />
</interface>
....
....
```

to this:

```
....
....
<interface type='network'>
  <source network='pvtnet01' />
  <model type='virtio' />
</interface>
....
....
```

and upon saving the change will be confirmed. Looking at the file again it can be seen that the *MAC address* and *address type* are automatically reconfigured.

Domain vm03 XML configuration edited

Note: For Microsoft Windows use *<model type='rtl8139' />* instead of *<model type='virtio' />*

Restart *vm03*.

```
virsh # start vm03
Domain vm03 started
```

9.2.1 IP configuration for guests

9.2.1.1 Debian guest

Configure an IP address on Debian guest *vm03* and *vm04*, edit each to have a similar look to the example below which reflects the file for *vm03*.

```
ada:~$ sudo vi /etc/network/interfaces
# This file describes the network interfaces available on your system
# and how to activate them. For more information, see interfaces(5).

source /etc/network/interfaces.d/*

# The loopback network interface
auto lo
iface lo inet loopback

auto eth0
iface eth0 inet static
    address 10.1.1.3
    netmask 255.255.255.0
    network 10.1.1.0
    broadcast 10.1.1.255
    gateway 10.1.1.1
    dns-nameservers 8.8.8.8

:wq!
```

Restart the *networking* service.

```
ada:~$ sudo /etc/init.d/networking restart
[ ok ] Restarting networking (via systemctl):
networking.service.
```

Confirm the IP settings.

```
ada:~$ ip addr list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default
    link/ether 52:54:00:6c:df:6c brd ff:ff:ff:ff:ff:ff
    inet 10.1.1.3/24 brd 10.1.1.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet6 fe80::5054:ff:fe6c:df6c/64 scope link
        valid_lft forever preferred_lft forever
```

Test between each guest once completed. First from *vm03* to *vm04*.

```
vm03:~$ ping -c1 10.1.1.4
PING 10.1.1.3 (10.1.1.4) 56(84) bytes of data.
64 bytes from 10.1.1.4: icmp_seq=1 ttl=64 time=0.128 ms

--- 10.1.1.4 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.139/0.139/0.139/0.000 ms
```

and now from *vm04* to *vm03*.

```
vm04:~$ ping -c1 10.1.1.3
PING 10.1.1.3 (10.1.1.3) 56(84) bytes of data.
64 bytes from 10.1.1.3: icmp_seq=1 ttl=64 time=0.139 ms

--- 10.1.1.3 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.139/0.139/0.139/0.000 ms
```

9.2.1.2 Windows guest

Configure an IP address on the *vm02* guest.

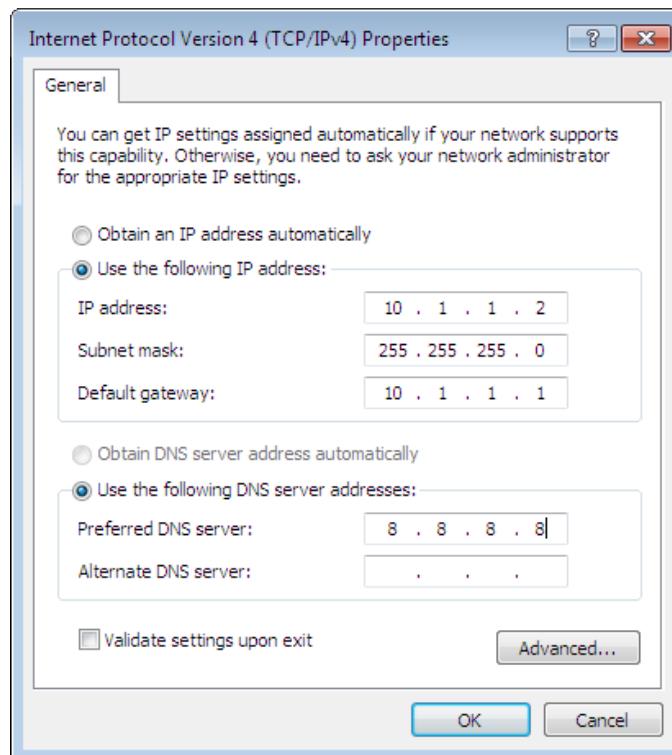


Illustration 11: Windows IPv4 address assignment

For testing purposes allow *ICMPv4 ECHO_REQ* on the Windows guest.

```
Control Panel --> System and security --> Windows Firewall  
Advanced settings --> Inbound rules --> New rule --> Custom rule
```

Protocol and ports: Protocol: **ICMPv4**

```
--> Customize,  
--> Choose Specific ICMP types, check the box echo request.
```

Open a command prompt and ping the domains *vm03* and *vm04* just configured.

```
C:> ping -n 1 10.1.1.3
```

```
Pinging 10.1.1.3 with 32 bytes of data:  
Reply from 10.1.1.3: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 10.1.1.3:  
Pakets: Sent = 1, Received = 1, Lost = 0 (0% loss),  
Approximate round trip times in milli-seconds:  
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

```
C:> ping -n 1 10.1.1.4
```

```
Pinging 10.1.1.4 with 32 bytes of data:  
Reply from 10.1.1.4: bytes=32 time<1ms TTL=64
```

```
Ping statistics for 10.1.1.4:  
Pakets: Sent = 1, Received = 1, Lost = 0 (0% loss),  
Approximate round trip times in milli-seconds:  
Minimum = 0ms, Maximum = 0ms, Average = 0ms
```

9.3 Configure the router *vm01-Debian_8*

9.3.1 Establish forwarding

By default the Debian 8 guest acting as a router *vm01* has IP forwarding disabled.

```
vm01:~$ cat /proc/sys/net/ipv4/ip_forward  
0
```

Enable forwarding by switching the value to 1.

```
vm01:~$ sudo bash -c 'echo 1 > /proc/sys/net/ipv4/ip_forward'
```

This interface configuration is similar to the guests except that two interfaces are required; one for the *Internet* and the other to the *pvtNet01*.

```
virsh # shutdown vm01  
Domain vm01 is being shutdown  
  
virsh # domstate vm01  
shut off
```

```
virsh # edit vm01
```

Change:

```
....  
<interface type='bridge'>  
  <mac address='52:54:00:c5:50:05' />  
  <source bridge='br0' />  
  <model type='virtio' />  
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />  
</interface>  
....
```

to this:

```
....  
<interface type='bridge'>  
  <mac address='52:54:00:c5:50:05' />  
  <source bridge='br0' />  
  <model type='virtio' />  
  <address type='pci' domain='0x0000' bus='0x00' slot='0x03' function='0x0' />  
</interface>  
<interface type='network'>  
  <source network='pvtnet01' />  
  <model type='virtio' />  
</interface>  
....
```

Domain vm01 XML configuration edited.

```
virsh # start vm01
```

9.3.2 Configure the interfaces on the router guest

Edit the */etc/network/interfaces* file as follows:

```
vm01:~$ sudo vi /etc/network/interfaces  
# This file describes the network interfaces available on your system  
# and how to activate them. For more information, see interfaces(5).  
  
source /etc/network/interfaces.d/*  
  
# The loopback network interface  
auto lo  
iface lo inet loopback  
  
# The primary interface eth0  
auto eth0  
iface eth0 inet dhcp  
  
# The secondary interface eth1 connected to pr  
auto eth1  
iface eth1 static  
  address 10.1.1.1  
  netmask 255.255.255.0  
  network 10.1.1.0  
  broadcast 10.1.1.255  
  dns-nameservers 8.8.8.8  
  
:wq!
```

Restart the *networking* service.

```
vm01:~$ sudo /etc/init.d/networking restart
[ ok ] Restarting networking (via systemctl): networking.service.
```

Confirm the IP settings.

```
vm01:~$ ip addr list
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
        inet 127.0.0.1/8 scope host lo
            valid_lft forever preferred_lft forever
        inet6 ::1/128 scope host
            valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default
    link/ether 52:54:00:fe:26:b7 brd ff:ff:ff:ff:ff:ff
        inet 213.74.34.85/24 brd 213.74.34.255 scope global dynamic eth0
            valid_lft 429sec preferred_lft 429sec
        inet6 fe80::5054:ff:fe26:b7/64 scope link
            valid_lft forever preferred_lft forever
3: eth1: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default
    link/ether 52:54:00:3c:a8:87 brd ff:ff:ff:ff:ff:ff
        inet 10.1.1.2/24 brd 10.1.1.255 scope global eth1
            valid_lft forever preferred_lft forever
        inet6 fe80::5054:ff:fe3c:a887/64 scope link
            valid_lft forever preferred_lft forever
```

Inspect the routes table.

```
vm01:~$ ip route list
default via 213.74.34.1 dev eth0
default via 213.74.34.1 dev eth0 proto static metric 1024
10.1.1.0/24 dev eth1 proto kernel scope link src 10.1.1.1
213.74.34.0/24 dev eth0 proto kernel scope link src 213.74.34.85
```

Ping hosts on the *pvteth0*.

```
vm01:~$ ping -c1 10.1.1.2
PING 10.1.1.2 (10.1.1.2) 56(84) bytes of data.
64 bytes from 10.1.1.2: icmp_seq=1 ttl=128 time=0.502 ms

--- 10.1.1.2 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.502/0.502/0.502/0.000 ms

vm01:~$ ping -c1 10.1.1.3
PING 10.1.1.3 (10.1.1.3) 56(84) bytes of data.
64 bytes from 10.1.1.3: icmp_seq=1 ttl=64 time=0.264 ms

--- 10.1.1.3 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.264/0.264/0.264/0.000 ms
```

```
vm01:~$ ping -c1 10.1.1.4
PING 10.1.1.4 (10.1.1.4) 56(84) bytes of data.
64 bytes from 10.1.1.4: icmp_seq=1 ttl=64 time=0.169 ms

--- 10.1.1.4 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 0.169/0.169/0.169/0.000 ms
```

Ping a public IP address.

```
vm01:~$ ping -c1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=56 time=28 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 26.562/26.562/26.562/0.000 ms
```

9.3.3 Configure IP Masquerade on the router

The IP Masquerade is a resource used so that a set of machines may use a single IP address. This permits the hidden nodes on a private network, such as 10.1.1.0/24 to access the public network; but they cannot directly accept external connections; only through the machine that has the real IP. Traffic from the private network to the Internet must have the private source IP address replaced with the Masquerade public IP address. Outward connections must be tracked so incoming returning traffic can be correctly identified and the correct private IP address swapped in the packet header for the public IP address before forwarding to the private network. This is achievable because of a GNU/Linux feature called Connection Tracking (conntrack). While on the public network the source IP address is masqueraded as if it came from the GNU/Linux server.

```
vm01:~$ sudo iptables --table nat \
--append POSTROUTING \
--source 10.1.1.0/24 \
--out-interface eth0 \
--jump MASQUERADE
```

9.4 Review the bridge control on the host

Have a look at what has happened on the host.

```
vm01:~$ sudo brctl show
bridge name      bridge id          STP enabled    interfaces
br0              8000.305a3a083921    no             eth0
                                         vnet0
                                         vnet1
                                         vnet5
                                         pvtbr0-nic
                                         vnet2
                                         vnet3
                                         vnet4
                                         vnet6

pvtbr0           8000.5254002db430    yes            pvtbr0-nic
                                         vnet2
                                         vnet3
                                         vnet4
                                         vnet6
```

9.5 Confirm that the guests have Internet access

Confirm each of the domain guests have Internet access through *vm01*.

9.5.1 vm02

```
C:> ping -n 1 8.8.8.8

Pinging 8.8.8.8 with 32 bytes of data:
Reply from 8.8.8.8: bytes=32 time<1ms TTL=64

Ping statistics for 8.8.8.8:
Pakets: Sent = 1, Received = 1, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
Minimum = 25ms, Maximum = 25ms, Average = 25ms
```

9.5.2 vm03

```
vm3:~$ ping -c1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=55 time=18.8 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 18.844/18.844/18.844/0.000 ms
```

9.5.3 vm04

```
vm4:~$ ping -c1 8.8.8.8
PING 8.8.8.8 (8.8.8.8) 56(84) bytes of data.
64 bytes from 8.8.8.8: icmp_seq=1 ttl=55 time=25.2 ms

--- 8.8.8.8 ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0ms
rtt min/avg/max/mdev = 25.204/25.204/25.204/0.000 ms
```

10. Graphical management

Now that the basics are mastered it is time to check out the KVM graphical manager.

```
ada:~$ virt-manager
```

This can be run on the hypervisor or remotely where it will connect to the hypervisor using *Secure SHell (SSH)*. When ran initially on a machine that is not the hypervisor it shows *localhost (QEMU) - Not Connected*. This simply demonstrates that there is no hypervisor on the local machine. To connect to the hypervisor select *File --> Add Connection....*

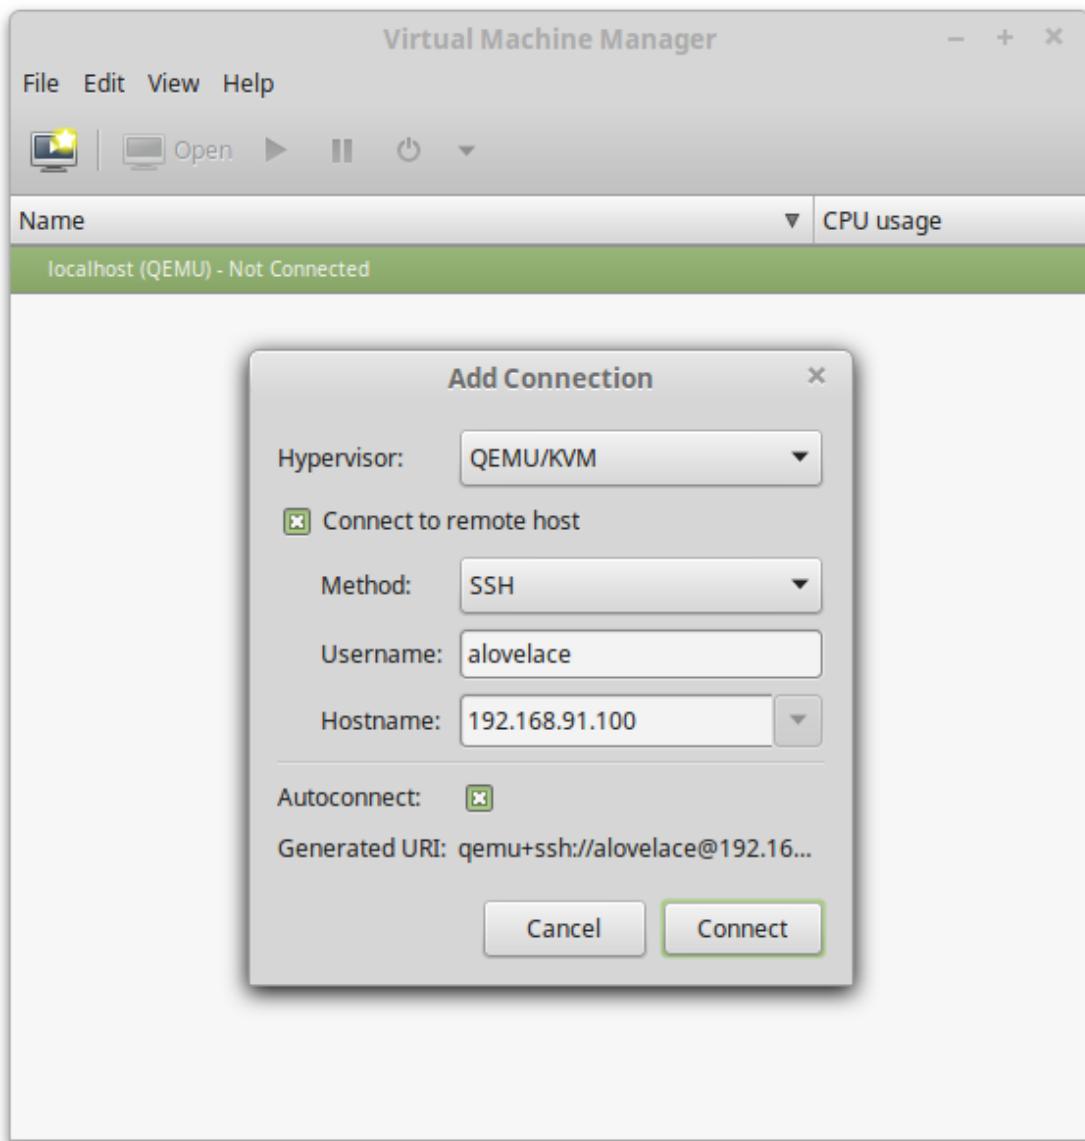


Illustration 12: Hypervisor 'Add connection' screen

Once logged in all of the domains are displayed in the initial window with the CPU usage graph for each.

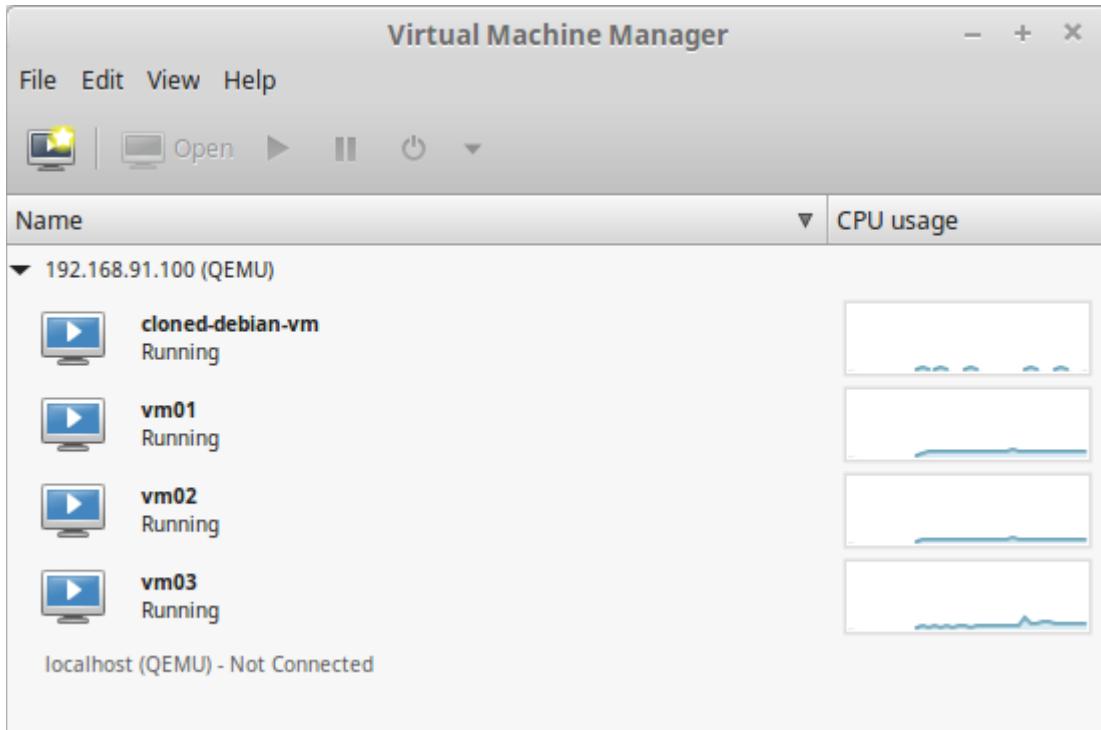


Illustration 13: Domain dashboard view

For each it is possible with a right-mouse click for each individual domain to:

- Pause
- Shutdown
- Migrate
- Open.

Select *Open* or double click on the domain to run a viewer to the domain VM in a similar fashion to that with *virt-viewer*. If the *View --> Details* is selected it is possible to adjust the domain options which will become active after the next domain reboot.

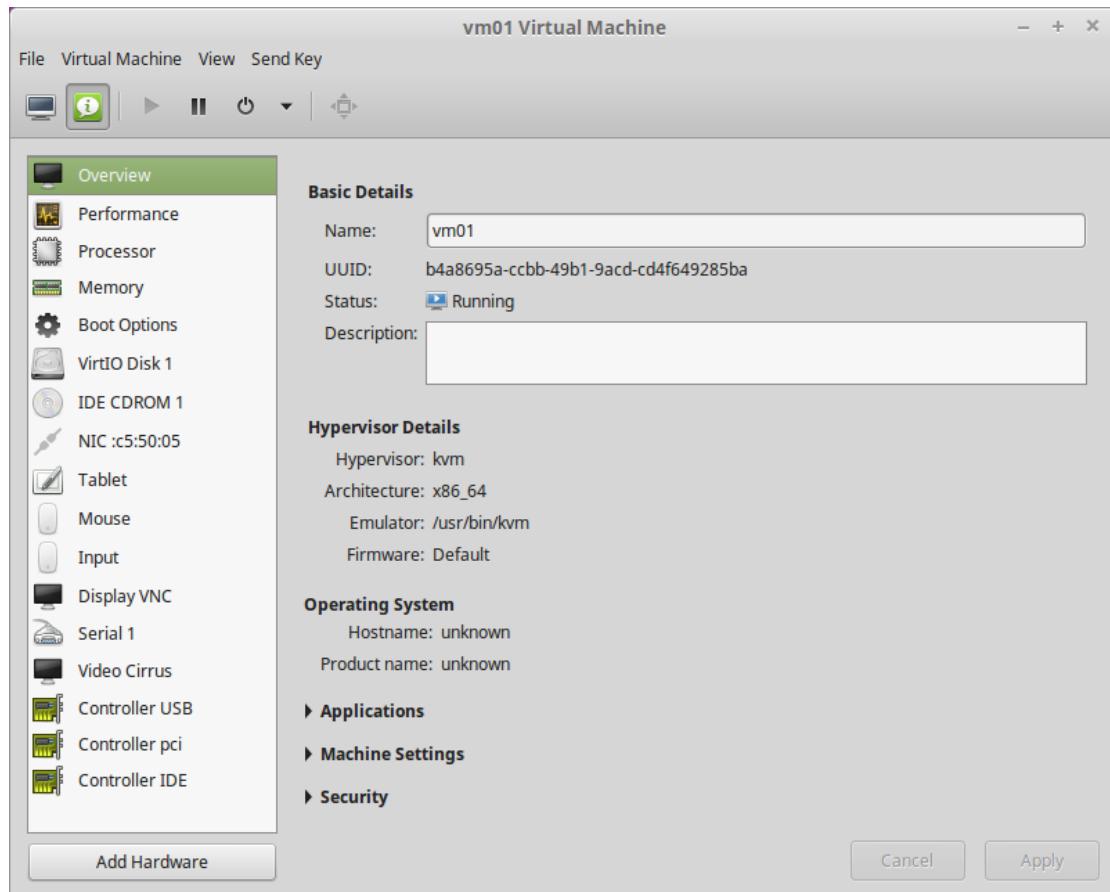


Illustration 14: Domain details

This page is intentionally blank

11. Importing image from other virtualisation platform

What if an image is received from another platform. For example an Open Virtualisation Archive (OVA) file from Oracle VirtualBox. Well the OVA file is actually a tar archive and can be extracted to show the Virtual Machine Disk (.vmdk) file within.

For KVM it is better to use the QCOW2 format.

Extract the files from the .ova file.

```
ada:~$ tar -xvf vm.ova
vm.ovf
vm.vmdk

ada:~$ qemu-img convert -O qcows2 vm-disk1.vmdk vm-disk1.qcows2
vm-disk1.qcows2    vm-disk1.vmdk    vm.ova    vm.ovf

ada:~$ cp vm-disk1.qcows2 /virt/kvm/images
```

If the details of the original VM are not readily available like CPU, RAM etc. they can be found in the Open Virtualisation Format (OVF) description file vm-vbox.ovf.

11.1 Install the KVM Domain with the import file

The existing disk that was created on another virtualisation platform can be imported when building the KVM domain. Build the KVM domain with the existing disk indicated in the disk path. The comma delimited options indicate the device is of the type disk and the disk device driver is *virtio*.

```
ada:~$ virt-install \
--connect qemu:///system \
--virt-type=kvm \
--name vm-OVA \
--ram 2048 \
--vcpus=2 \
--disk path=/virt/kvm/images/vm-disk1.qcows2,device=disk,bus=virtio \
--graphics vnc,listen=0.0.0.0 \
--noautoconsole \
--os-type linux \
--os-variant debian8 \
--network=bridge:br0 \
--import

Starting install...
Creating domain... | 0 B 00:01
Connected to domain vm-OVA

Escape character is ^]

Domain creation completed. You can restart your domain by running:
virsh --connect qemu:///system start vm-OVA
```

This page is intentionally blank

12. Appendix 1

This is a simplified minimal install process for Ubuntu Server.

12.1 Check if KVM is supported

If KVM is supported this will return a non-0 value.

```
ada:~$ egrep -c '(vmx|svm)' /proc/cpuinfo
4
```

12.2 Install the KVM elements

Install the required elements.

```
ada:~$ sudo apt-get install bridge-utils
ada:~$ sudo apt-get install qemu-kvm libvirt-bin ubuntu-vm-builder
ada:~$ sudo apt-get install qemu-utils
ada:~$ sudo apt-get install virt-manager
ada:~$ sudo apt-get install libspice-client-gtk-2.0-4
ada:~$ sudo apt-get install python-spice-client-gtk
ada:~$ sudo apt-get install virt-viewer
```

12.3 Add user to the libvirt daemon group

Add the user to the *libvird* group.

```
ada:~$ sudo adduser `id -un` libvird
ada:~$ cat /etc/group | grep libvird
libvird:x:133:dobriain
```

12.4 VM network

Note the bridge *virbr0* created by KVM.

```
ada:~$ sudo brctl show
[sudo] password for dobriain:
bridge name      bridge id           STP enabled     interfaces
virbr0          8000.000000000000       yes
```

IP network assigned to the bridge.

```
ada:~$ ip addr show dev virbr0
5: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue
    state DOWN group default qlen 1000
    link/ether 00:00:00:00:00:00 brd ff:ff:ff:ff:ff:ff
    inet 192.168.122.1/24 brd 192.168.122.255 scope global virbr0
        valid_lft forever preferred_lft forever

ada:~$ sudo brctl show
[sudo] password for dobriain:
bridge name      bridge id      STP enabled      interfaces
virbr0          8000.fe540011fa18  yes            vnet0
                                         vnet1
```

12.5 Install a VM

In the *Virtual Machine Manager*, select:

- File - New Virtual Machine

Under *Create a new virtual machine* select *local install media*.

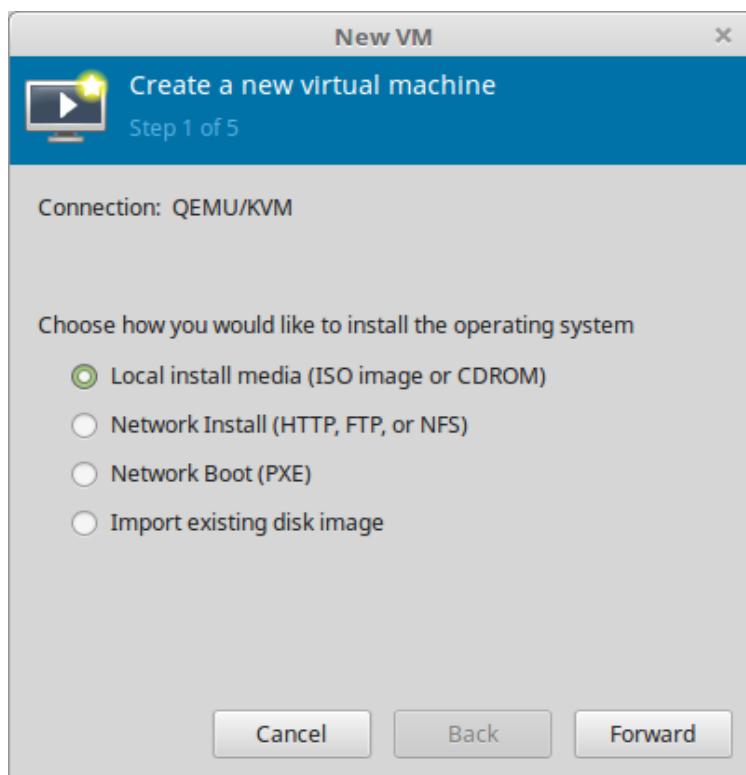


Illustration 15: Create VM

Then define the ISO image or the DVD/CD.

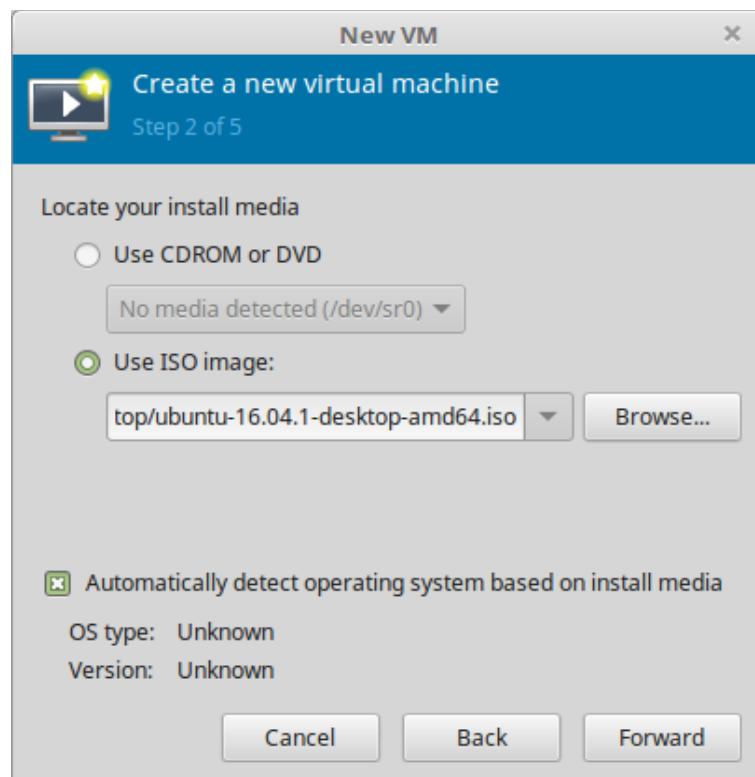


Illustration 16: Select image

Specify Memory and CPU details for the VM.

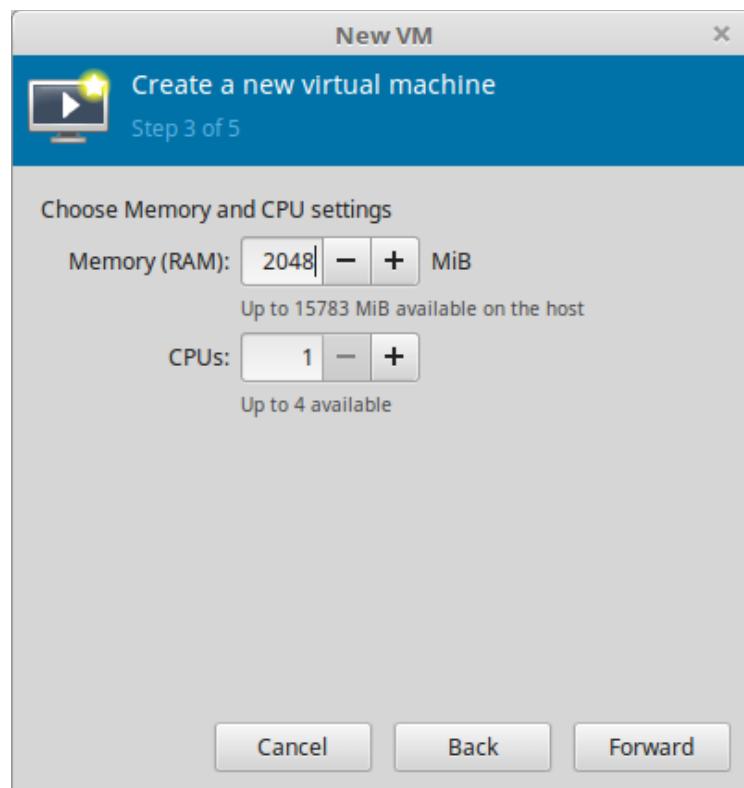


Illustration 17: Memory and CPU

Define the size of the disk image.

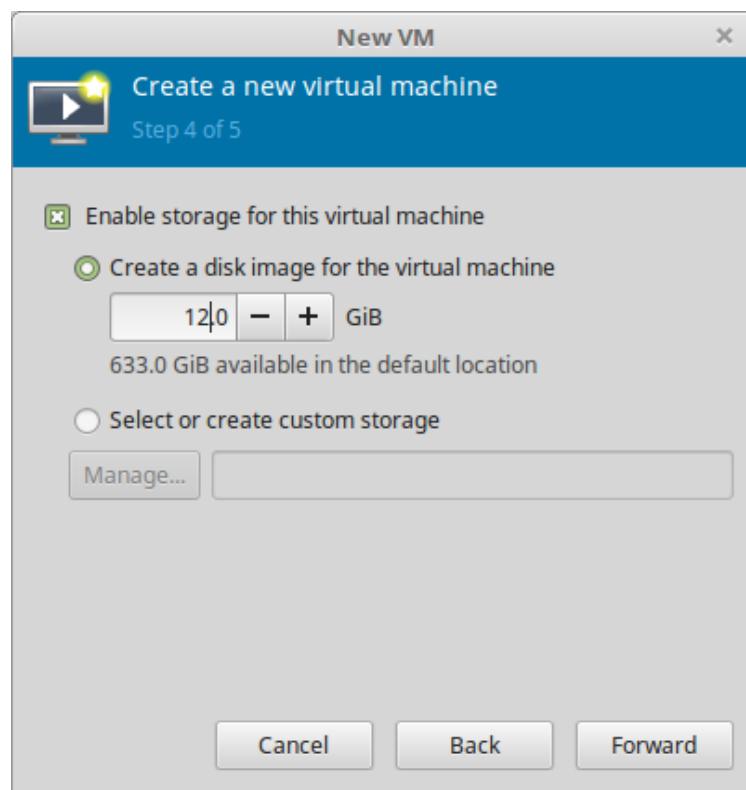


Illustration 18: Create disk image

Give the VM a name and make sure the VM is in the *default* network.

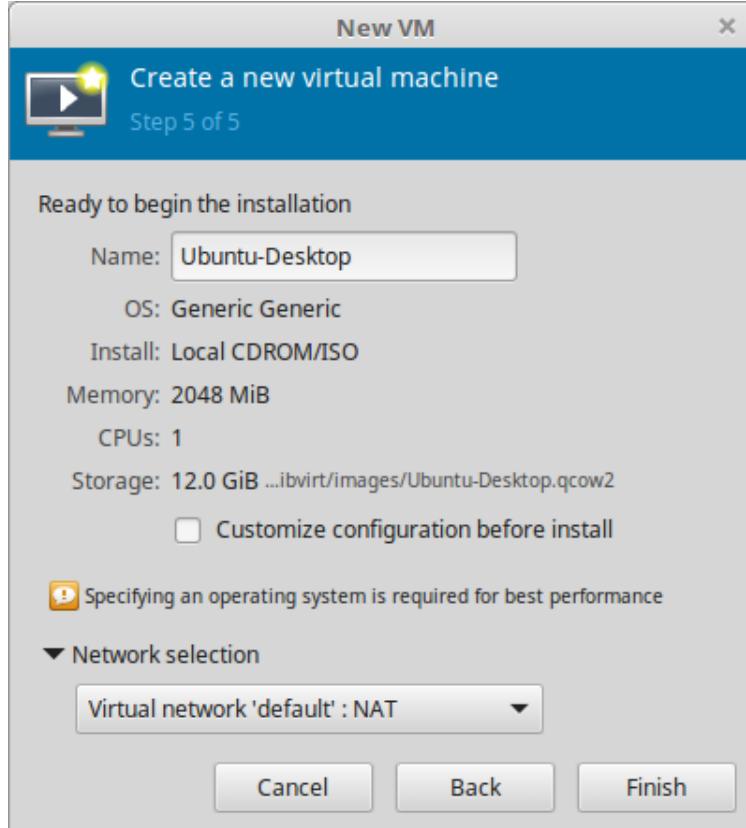


Illustration 19: Confirm configuration and finish

12.6 Accessing the VM's

Power on the VM by clicking on the play button on the Virtual Machine Manager. Use the *virt-viewer* utility to access VMs. By default the VMs are set to use spice.

```
$ virt-viewer Ubuntu-Desktop
```

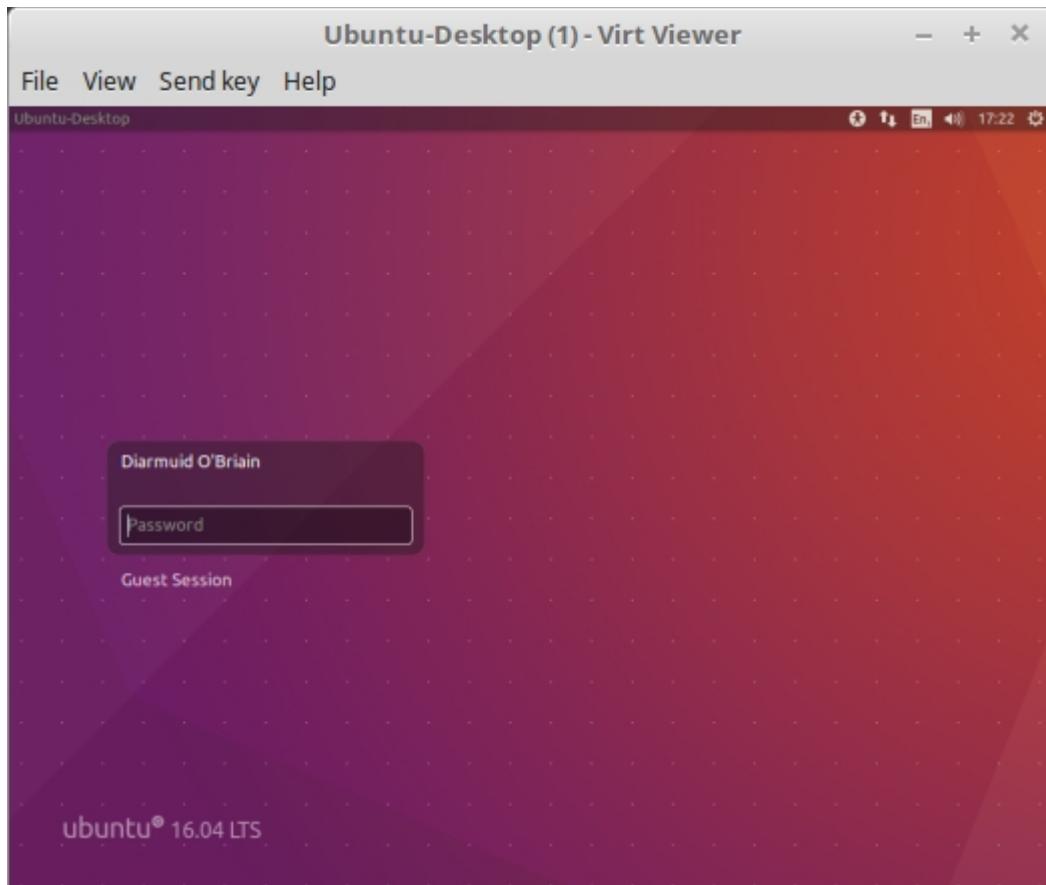


Illustration 20: Ubuntu Desktop VM

```
$ virt-viewer Ubuntu-Server
```

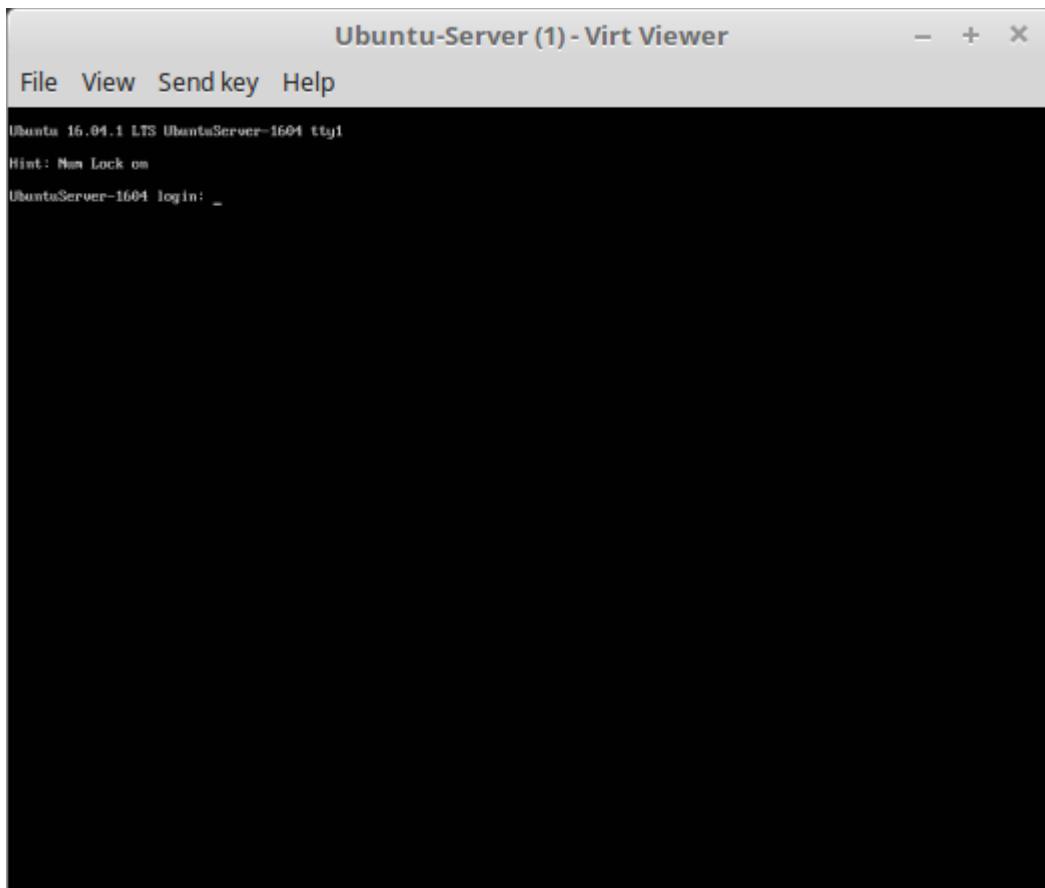


Illustration 21: Ubuntu Server VM

12.7 IP address on KVM Domains

Access the KVM shell.

```
$ virsh
```

List the Virtual Machines, called *domains* in KVM.

```
virsh # list
 Id      Name           State
 -----
 2      Ubuntu-Server   running
 3      Ubuntu-Desktop  running
```

Using the *virtsh* get the IP address of the Virtual Machines.

```
virsh # domifaddr Ubuntu-Desktop
 Name      MAC address      Protocol      Address
 -----
 vnet0    52:54:00:76:be:a3  ipv4        192.168.122.117/24

virsh # domifaddr Ubuntu-Server
 Name      MAC address      Protocol      Address
 -----
 vnet1    52:54:00:11:fa:18  ipv4        192.168.122.125/24
```

Confirm IP addresses directly on domain.

```
ubuntu@Ubuntu-Desktop:~$ ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 52:54:00:76:be:a3 brd ff:ff:ff:ff:ff:ff
    inet 192.168.122.117/24 brd 192.168.122.255 scope global dynamic ens3
        valid_lft 3305sec preferred_lft 3305sec
    inet6 fe80::830a:1a58:b3b9:a1da/64 scope link
        valid_lft forever preferred_lft forever

ubuntu@UbuntuServer:~$ ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: ens3: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
    link/ether 52:54:00:11:fa:18 brd ff:ff:ff:ff:ff:ff
    inet 192.168.122.125/24 brd 192.168.122.255 scope global ens3
        valid_lft forever preferred_lft forever
    inet6 fe80::5054:ff:fe11:fa18/64 scope link
        valid_lft forever preferred_lft forever
```

This page is intentionally blank

13. Abbreviations

BIOS	Basic Input/Output System
CPU	Central Processing Unit
conntrack	Connection Tracking
DHCP	Dynamic Host Configuration Protocol
XML	eXtensible Markup Language
GB	Gigabytes
GNU	GNU's not Unix
HVM	Hardware-assisted Virtual Machine
I/O	Input/Output
ISO	International Organization for Standardisation 9660 image
IP	Internet Protocol
KVM	Kernel-based Virtual Machine
virsh	libvirt based command line interface tool for managing guests and the hypervisor.
LXC	LinuX Container
LM	Long Mode
MAC	Medium Access Control
MB	Megabytes
NAT	Network Address Translation
OVA	Open Virtualisation Archive
OVF	Open Virtualisation Format
OS	Operating System
PV	Para Virtualisation
QCOW	QEMU Copy On Write
QCOW2	QEMU Copy On Write version 2
QCOW3	QEMU Copy On Write version 3
QEMU	Quick EMULATOR
RAM	Random Access Memory
ROI	Return on Investment
SSH	Secure SHell
SVM	Secure Virtual Machine
SPICE	Simple Protocol for Independent Computing Environments
libvirt	toolkit to manage virtualisation hosts
UUID	Universally Unique IDentifier

VIM	Vi IMproved
VMX	Virtual Machine eXtensions
VM	Virtual Machines
VNC	Virtual Network Computing
VT-x	Virtualisation Technology - x86 architectures
VT-d	Virtualisation Technology for Directed I/O
ESX	VMware vSphere

14. Bibliography

KVM. (2017). KVM virtualization solution for Linux Website [online]. Available at: <http://www.linux-kvm.org> [Accessed: 4 Jan 2017]

Libvirt. (2017). Libvirt Virtualisation API Website [online]. Available at: <https://libvirt.org> [Accessed: 4 Jan 2017]

QEMU. (2017). QEMU Emulator Website [online]. Available at: <http://wiki.qemu.org> [Accessed: 3 Jan 2017]

This page is intentionally blank