



# Red Hat Enterprise Linux

## 6

# Virtualization Getting Started Guide

---

An introduction to virtualization concepts

Jiri Herrmann  
Laura Novich

Yehuda Zimmerman  
Jacquelynn East

Dayle Parker  
Scott Radvan



# Red Hat Enterprise Linux 6 Virtualization Getting Started Guide

---

## An introduction to virtualization concepts

Jiri Herrmann  
Red Hat Customer Content Services  
jherrman@redhat.com

Yehuda Zimmerman  
Red Hat Customer Content Services  
yzimmerm@redhat.com

Dayle Parker  
Red Hat Customer Content Services

Laura Novich  
Red Hat Customer Content Services

Jacquelynn East  
Red Hat Customer Content Services

Scott Radvan  
Red Hat Customer Content Services

## Legal Notice

Copyright © 2017 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

## Abstract

The Red Hat Enterprise Linux Virtualization Getting Started Guide describes the basics of virtualization and the virtualization products and technologies that are available with Red Hat Enterprise Linux.

## Table of Contents

<b>Chapter 1. General Introduction to Virtualization</b> .....	<b>2</b>
1.1. What is Virtualization?	2
1.2. Virtualization Solutions	2
<b>Chapter 2. Why Use Virtualization?</b> .....	<b>4</b>
2.1. Virtualization Costs	4
2.2. Performance	4
2.3. Migration	5
2.3.1. Benefits of Migrating Virtual Machines	5
2.3.2. Virtualized to Virtualized Migration (V2V)	6
2.4. Security	6
2.4.1. Virtualization Security Features	6
2.5. Disaster Recovery	7
<b>Chapter 3. Introduction to Red Hat Virtualization Products and Features</b> .....	<b>8</b>
3.1. KVM and Virtualization in Red Hat Enterprise Linux	8
3.2. libvirt and libvirt Tools	10
3.3. Virtualized Hardware Devices	10
3.3.1. Virtualized and Emulated devices	11
3.3.2. Paravirtualized Devices	13
3.3.3. Physical Host Devices	14
3.4. Guest CPU Models	15
3.5. Storage	16
3.5.1. Storage Pools	16
3.5.2. Storage Volumes	16
3.5.3. Emulated Storage Devices	16
3.5.4. Host Storage	17
3.6. Virtual Networking	18
<b>Chapter 4. Quick Start Tutorial on Virtualization in Red Hat Enterprise Linux 6</b> .....	<b>20</b>
4.1. Tutorial overview	20
4.2. Basic Requirements and Setup	20
4.3. Creating a Virtual Machine with Virtual Machine Manager	21
4.3.1. Introduction to Virtual Machine Manager	21
4.3.2. Creating a Virtual Machine with Virtual Machine Manager	22
4.3.3. Exploring the Guest Virtual Machine	27
<b>Chapter 5. Virtualization Tools</b> .....	<b>30</b>
5.1. virsh	30
5.2. virt-manager	30
5.3. virt-install	30
5.4. guestfish	30
5.5. Other Useful Tools	31
<b>Appendix A. Revision History</b> .....	<b>36</b>

# Chapter 1. General Introduction to Virtualization

## 1.1. What is Virtualization?

*Virtualization* is a broad computing term used for running software, usually multiple operating systems, concurrently and in isolation from other programs on a single system. Virtualization is accomplished by using a *hypervisor*. This is a software layer or subsystem that controls hardware and enables running multiple operating systems, called *virtual machines (VMs)* or *guests*, on a single (usually physical) machine. This machine with its operating system is called a *host*. There are several virtualization methods:

### Full virtualization

Full virtualization uses an unmodified version of the guest operating system. The guest addresses the host's CPU using a channel created by the hypervisor. Because the guest communicates directly with the CPU, this is the fastest virtualization method.

### Paravirtualization

Paravirtualization uses a modified guest operating system. The guest communicates with the hypervisor. The hypervisor passes the unmodified calls from the guest to the CPU and other interfaces, both real and virtual. Because the calls are routed through the hypervisor, this method is slower than full virtualization.

### Software virtualization (or emulation)

Software virtualization uses binary translation and other emulation techniques to run unmodified operating systems. The hypervisor translates the guest calls to a format that can be used by the host system. Because all calls are translated, this method is slower than virtualization. Note that Red Hat does not support software virtualization on Red Hat Enterprise Linux.

For information about the benefits of using virtualization, see [Chapter 2, Why Use Virtualization?](#)

## 1.2. Virtualization Solutions

Red Hat offers two major virtualization solutions with different user focuses and features:

### Red Hat Enterprise Linux

The ability to create, run, and manage virtual machines, as well as [a number of virtualization tools and features](#), are included in Red Hat Enterprise Linux 6. This solution supports a limited number of running guests per host, as well as a limited range of guest types. As such, virtualization on Red Hat Enterprise Linux can be useful for example to developers who require testing in multiple environments, or to small businesses running several servers that do not have strict uptime requirements or service-level agreements (SLAs).

### Red Hat Virtualization

Red Hat Virtualization is based on the Kernel-based Virtual Machine ([KVM](#)) technology like virtualization on Red Hat Enterprise Linux is, but offers an enhanced array of features. Designed for enterprise-class scalability and performance, it enables management of your entire virtual infrastructure, including hosts, virtual machines, networks, storage, and users from a centralized graphical interface.

Red Hat Virtualization can be used by enterprises running larger deployments or mission-critical applications. Examples of large deployments suited to Red Hat Virtualization include databases, trading platforms, and messaging systems that must run continuously without any downtime.



## Note

This guide provides information about virtualization on Red Hat Enterprise Linux. For more information about Red Hat Virtualization, or to download a fully supported 60-day evaluation version, see <http://www.redhat.com/en/technologies/virtualization/enterprise-virtualization>

## Chapter 2. Why Use Virtualization?

Virtualization can be useful both for server deployments and individual desktop stations. Desktop virtualization offers cost-efficient centralized management and better disaster recovery. In addition, by using connection tools such as `ssh`, it is possible to connect to a desktop remotely.

When used for servers, virtualization can benefit not only larger networks, but also deployments with more than a single server. Virtualization provides live migration, high availability, fault tolerance, and streamlined backups.

### 2.1. Virtualization Costs

Virtualization can be expensive to introduce, but it often saves money in the long term. Consider the following benefits:

#### Less power

Using virtualization negates much of the need for multiple physical platforms. This equates to less power being drawn for machine operation and cooling, resulting in reduced energy costs. The initial cost of purchasing multiple physical platforms, combined with the machines' power consumption and required cooling, is drastically cut by using virtualization.

#### Less maintenance

Provided that adequate planning is performed before migrating physical systems to virtualized ones, less time is spent maintaining them. This means less money needs to be spent on parts and labor.

#### Extended life for installed software

Older versions of software may not be able to run directly on more recent physical machines. By running older software virtually on a larger, faster system, the life of the software may be extended while taking advantage of better performance from a newer system.

#### Predictable costs

A Red Hat Enterprise Linux subscription provides support for virtualization at a fixed rate, making it easy to predict costs.

#### Less space

Consolidating servers onto fewer machines means less physical space is required for computer systems.

### 2.2. Performance

Older virtualization versions supported only a single CPU. As a result, virtual machines experienced noticeable performance limitations. This created a long-lasting misconception that virtualization solutions are slow.

This is no longer the case. Modern virtualization technology has greatly improved the speed of virtual machines. Benchmarks show that virtual machines can run typical server applications nearly as efficiently as bare-metal systems:

- Red Hat Enterprise Linux 6.4 and KVM recorded [an industry-leading TPC-C benchmark](#) with an IBM DB2 database running in an entirely virtualized x86 environment and delivering 88% of bare-metal performance. Due to resource demands, databases have previously been reserved for bare-metal



performance. Due to resource demands, databases have previously been reserved for bare metal deployments only.

- ✦ The industry standard SAP Sales and Distribution (SD) Standard Application Benchmark found that Red Hat Enterprise Linux 6.2 and KVM [performs at the virtualization efficiency of 85%](#) compared to a bare-metal system running on identical hardware.
- ✦ Red Hat Enterprise Linux 6.1 and KVM achieved [record-setting virtualization performance in the SPECvirt\\_sc2010 benchmark](#) recorded by the Standard Performance Evaluation Corporation (SPEC), setting the best virtual performance mark of any published SPECvirt result. The SPECvirt\_sc2010 metric measures the end-to-end performance of system components in virtualized data center servers.



## Note

For more information on performance tuning for virtualization, refer to the [Red Hat Enterprise Linux 6 Virtualization Tuning and Optimization Guide](#).

## 2.3. Migration

*Migration* describes the process of moving a guest virtual machine from one host to another. This is possible because the virtual machines are running in a virtualized environment instead of directly on the hardware. There are two ways to migrate a virtual machine: live and offline.

### Migration Types

#### Offline migration

An offline migration suspends the guest virtual machine, and then moves an image of the virtual machine's memory to the destination host. The virtual machine is then resumed on the destination host and the memory used by the virtual machine on the source host is freed.

#### Live migration

Live migration is the process of migrating an active virtual machine from one physical host to another. Note that this is not possible between all Red Hat Enterprise Linux releases. Consult the [Virtualization Administration Guide](#) for details.

### 2.3.1. Benefits of Migrating Virtual Machines

Migration is useful for:

#### Load balancing

When a host machine is overloaded, one or more of its virtual machines could be migrated to other hosts using live migration. Similarly, machines that are not running and tend to overload can be migrated using offline migration.

#### Upgrading or making changes to the host

When the need arises to upgrade, add, or remove hardware devices on a host, virtual machines can be safely relocated to other hosts. This means that guests do not experience any downtime due to changes that are made to hosts.

#### Energy saving

Virtual machines can be redistributed to other hosts and the unloaded host systems can be powered off to save energy and cut costs in low usage periods.

### Geographic migration

Virtual machines can be moved to other physical locations for lower latency or for other reasons.

When the migration process moves a virtual machine's memory, from Red Hat Enterprise Linux 6.3, the disk volume associated with the virtual machine is also migrated. This process is performed using live block migration.

Shared, networked storage can be used to store guest images to be migrated. When migrating virtual machines, it is recommended to use **libvirt**-managed storage pools for shared storage.



#### Note

For more information on migration, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

### 2.3.2. Virtualized to Virtualized Migration (V2V)

Red Hat Enterprise Linux 6 provides tools for converting virtual machines from other types of hypervisors to KVM. The **virt-v2v** tool converts and imports virtual machines from Xen, other versions of KVM, VMware ESX, and Hyper-V.

In Red Hat Enterprise Linux 6.5 and later, the **virt-v2v** tool converts and imports virtual machines with virtual disk image formats including VHDX, VMDK, OVF, raw, and qcow2.



#### Important

The V2V tool for Red Hat Enterprise Linux 6 has been deprecated, and it is recommended to use the Red Hat Enterprise Linux 7 version. For more information, refer to the [Red Hat Enterprise Linux 6 V2V Guide](#).

## 2.4. Security

A virtual machine uses SELinux and sVirt to improve security in virtualization. This section includes an overview of the security options available.

### 2.4.1. Virtualization Security Features

#### SELinux

Security-Enhanced Linux, or SELinux, provides Mandatory Access Control (MAC) for all Linux system, and thus benefits also Linux guests. Under the control of SELinux, all processes and files are given a *type*, and their access on the system is limited by fine-grained controls of various types. SELinux limits the abilities of an attacker and works to prevent many common security exploits such as buffer overflow attacks and privilege escalation.

SELinux strengthens the security model of Red Hat Enterprise Linux hosts and virtualized Red Hat Enterprise Linux guests. SELinux is configured and tested to work, by default, with all virtualization tools shipped with Red Hat Enterprise Linux 6.

### sVirt

sVirt is a technology included in Red Hat Enterprise Linux 6 that integrates SELinux and virtualization. It applies Mandatory Access Control (MAC) to improve security when using virtual machines, and improves security and hardens the system against hypervisor bugs that might be used to attack the host or another virtual machine.



#### Note

For more information on security for virtualization, refer to the [Red Hat Enterprise Linux 6 Virtualization Security Guide](#).

## 2.5. Disaster Recovery

Disaster recovery is quicker and easier when the systems are virtualized. On a physical system, if something serious goes wrong, a complete reinstall of the operating system is usually required, resulting in hours of recovery time. However, if the systems are virtualized this is much faster due to the migration ability. If the requirements for live migration are followed, virtual machines can be restarted on another host, and the longest possible delay would be in restoring guest data. Also, because each of the virtualized systems are completely separate to each other, one system's downtime will not affect any others.

## Chapter 3. Introduction to Red Hat Virtualization Products and Features

This chapter introduces the main virtualization products and features available in Red Hat Enterprise Linux.

### 3.1. KVM and Virtualization in Red Hat Enterprise Linux

KVM (Kernel-based Virtual Machine) is a full virtualization solution for Linux on AMD64 and Intel 64 hardware. It is built into the standard Red Hat Enterprise Linux 6 kernel and integrated with the Quick Emulator (QEMU), and it can run multiple, guest operating systems. The KVM hypervisor in Red Hat Enterprise Linux is managed with the `libvirt` API, and tools built for `libvirt` (such as `virt-manager` and `virsh`). Virtual machines are executed and run as multi-threaded Linux processes, controlled by these tools.

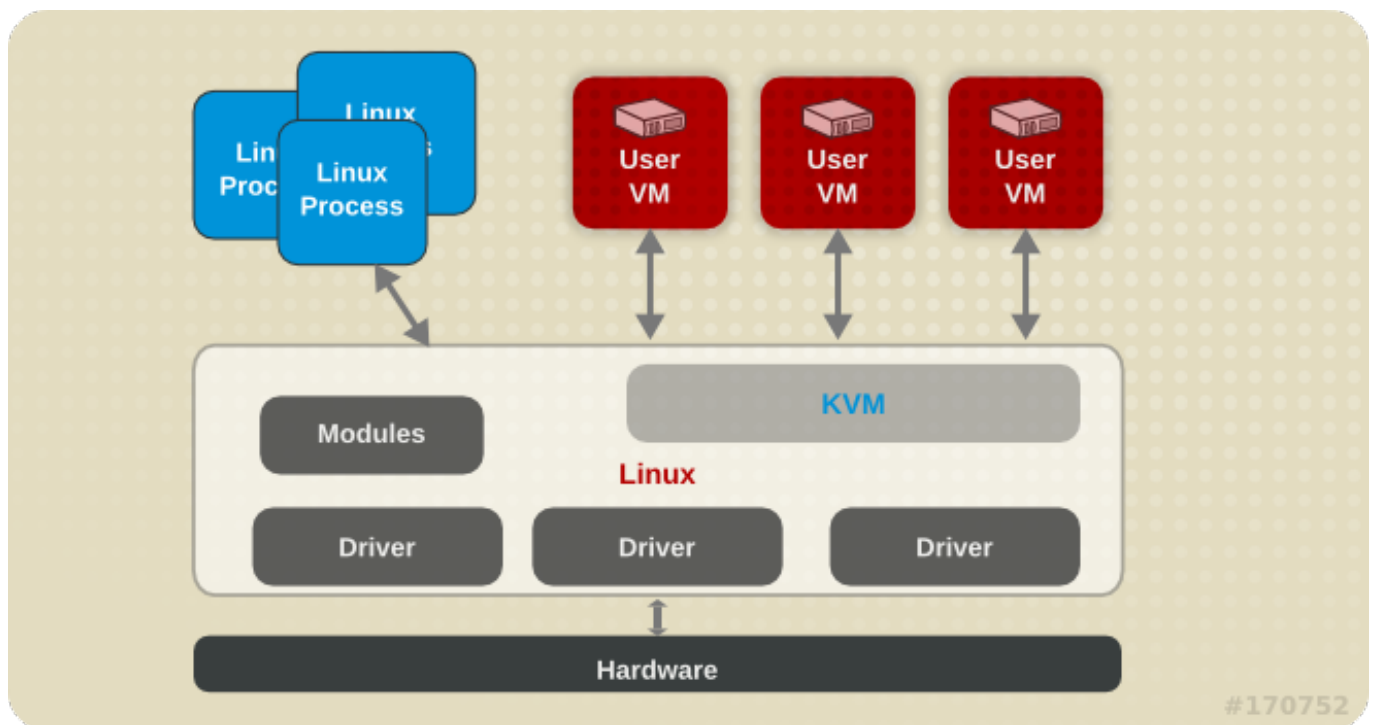


Figure 3.1. KVM architecture

Virtualization features supported by KVM on Red Hat Enterprise 6 include the following:

#### Overcommitting

The KVM hypervisor supports *overcommitting* of system resources. Overcommitting means allocating more virtualized CPUs or memory than the available resources on the system, so the resources can be dynamically swapped when required by one guest and not used by another. This can improve how efficiently guests use the resources of the host, and can make it possible for the user to require fewer hosts.



#### Important

Overcommitting involves possible risks to system stability. For more information on overcommitting with KVM, and the precautions that should be taken, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## KSM

*Kernel Same-page Merging (KSM)*, used by the KVM hypervisor, makes it possible for KVM guests to share identical memory pages. These shared pages are usually common libraries or other identical, high-use data. KSM allows for greater guest density of identical or similar guest operating systems by avoiding memory duplication.



### Note

For more information on KSM, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## QEMU guest agent

The *QEMU guest agent* runs on the guest operating system and enables the host machine to issue commands to the guest operating system.



### Note

For more information on the QEMU guest agent, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

## Hyper-V Enlightenment

KVM in Red Hat Enterprise Linux 6.5 and later implements several Hyper-V compatible functions that are used by Windows guests to improve performance and stability, enabling Windows guests to perform as if they were running on a Microsoft Hyper-V hypervisor.



### Note

For more information on Hyper-V Enlightenment, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

## KVM guest virtual machine compatibility

Red Hat Enterprise Linux 6 servers have certain support limits.

The following URLs explain the processor and memory amount limitations for Red Hat Enterprise Linux:

- ✦ For the host system: <https://access.redhat.com/site/articles/rhel-limits>
- ✦ For the KVM hypervisor: <https://access.redhat.com/site/articles/rhel-kvm-limits>

For a complete chart of supported operating systems and host and guest combinations refer to [Red Hat Customer Portal](#).



## Note

To verify whether your processor supports virtualization extensions and for information on enabling virtualization extensions if they are disabled, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## 3.2. libvirt and libvirt Tools

The *libvirt* package provides a hypervisor-independent virtualization API that can interact with the virtualization capabilities of a range of operating systems. It includes:

- ✦ A virtualization layer to securely manage virtual machines on a host.
- ✦ An interface for managing local and networked hosts.
- ✦ The APIs required to provision, create, modify, monitor, control, migrate, and stop virtual machines. Although multiple hosts may be accessed with **libvirt** simultaneously, the APIs are limited to single node operations.



## Note

Only operations supported by the hypervisor can be performed using libvirt.

**libvirt** focuses on managing single hosts and provides APIs to enumerate, monitor and use the resources available on the managed node, including CPUs, memory, storage, networking and Non-Uniform Memory Access (NUMA) partitions. The management tools do not need to be on the same physical machine as the machines on which the hosts are running. In such a scenario, the machine on which management tools run communicates with the machines on which the hosts are running using secure protocols.

Red Hat Enterprise Linux 6 supports **libvirt** and includes **libvirt**-based tools as its default method for virtualization management (as in Red Hat Virtualization Management).

The *libvirt* package is available as free software under the GNU Lesser General Public License. The *libvirt* project aims to provide a long term stable C API to virtualization management tools, running on top of varying hypervisor technologies. The *libvirt* package supports Xen on Red Hat Enterprise Linux 5, and it supports KVM on both Red Hat Enterprise Linux 5 and Red Hat Enterprise Linux 6.

Notably, libvirt also provides the two central tools for controlling virtualization on Red Hat Enterprise Linux 6: [virsh](#) and [virt-manager](#)

## 3.3. Virtualized Hardware Devices

Virtualization on Red Hat Enterprise Linux 6 presents three distinct types of system devices to virtual machines. The three types include:

- ✦ Virtualized and Emulated Devices
- ✦ Paravirtualized devices
- ✦ Physically shared devices

These hardware devices all appear as being physically attached to the virtual machine but the device drivers work in different ways.

### 3.3.1. Virtualized and Emulated devices

KVM implements many core devices for virtual machines as software. These emulated hardware devices are crucial for virtualizing operating systems. Emulated devices are virtual devices which exist entirely in software.

In addition, KVM provides emulated drivers. These form a translation layer between the virtual machine and the Linux kernel (which manages the source device). The device level instructions are completely translated by the KVM hypervisor. Any device of the same type (storage, network, keyboard, or mouse) that is recognized by the Linux kernel can be used as the backing source device for the emulated drivers.

#### Virtual CPUs (vCPUs)

A host system can have up to 160 virtual CPUs (vCPUs) that can be presented to guests for use, regardless of the number of host CPUs.

#### Emulated system components

The following core system components are emulated to provide basic system functions:

- Intel i440FX host PCI bridge
- PIIX3 PCI to ISA bridge
- PS/2 mouse and keyboard
- EvTouch USB Graphics Tablet
- PCI UHCI USB controller and a virtualized USB hub
- Emulated serial ports
- EHCI controller, virtualized USB storage and a USB mouse

#### Emulated storage drivers

Storage devices and storage pools can use these emulated devices to attach storage devices to virtual machines. The guest uses an emulated storage driver to access the storage pool.

Note that like all virtual devices, the storage drivers are not storage devices. The drivers are used to attach a backing storage device, file or storage pool volume to a virtual machine. The backing storage device can be any supported type of storage device, file, or storage pool volume.

##### The emulated IDE driver

KVM provides two emulated PCI IDE interfaces. An emulated IDE driver can be used to attach any combination of up to four virtualized IDE hard disks or virtualized IDE CD-ROM drives to each virtual machine. The emulated IDE driver is also used for virtualized CD-ROM and DVD-ROM drives.

##### The emulated floppy disk drive driver

The emulated floppy disk drive driver is used for creating virtualized floppy drives.

#### Emulated sound devices

Red Hat Enterprise Linux 6.1 and above provides an emulated (Intel) HDA sound device, **intel-hda**. This device is supported on the following guest operating systems:

- Red Hat Enterprise Linux 6, for the 32-bit AMD and Intel architecture, and AMD64 and Intel 64 architectures
- Red Hat Enterprise Linux 5, for i386, and the 32-bit AMD and Intel architecture and Intel 64 architectures
- Red Hat Enterprise Linux 4, for i386 and the 32-bit AMD and Intel architecture and Intel 64 architectures
- Windows 7, for i386 and AMD64 and Intel 64 architectures
- Windows 2008 R2, for the AMD64 and Intel 64 architecture



### Note

The following emulated sound devices are also available, but are not recommended due to compatibility issues with certain guest operating systems:

- **ac97**, an emulated Intel 82801AA AC97 Audio compatible sound card
- **es1370**, an emulated ENSONIQ AudioPCI ES1370 sound card

### Emulated graphics cards

The following emulated graphics devices are provided:

- A Cirrus CLGD 5446 PCI VGA card
- A standard VGA graphics card with Bochs VESA extensions (hardware level, including all non-standard modes)

Guests can connect to these devices with the Simple Protocol for Independent Computing Environments (SPICE) protocol or with the Virtual Network Computing (VNC) system.

### Emulated network devices

The following emulated network devices are provided:

- The **e1000** device emulates an Intel E1000 network adapter (Intel 82540EM, 82573L, 82544GC).
- The **rt18139** device emulates a Realtek 8139 network adapter.

### Emulated watchdog devices

Red Hat Enterprise Linux 6 provide two emulated watchdog devices. A watchdog can be used to automatically reboot a virtual machine when it becomes overloaded or unresponsive.

The *watchdog* package must be installed on the guest.

The two devices available are:

- **i6300esb**, an emulated Intel 6300 ESB PCI watchdog device. It is supported in guest operating system Red Hat Enterprise Linux versions 6.0 and above, and is the recommended device to use.



- **ib700**, an emulated iBase 700 ISA watchdog device. The **ib700** watchdog device is only supported in guests using Red Hat Enterprise Linux 6.2 and above.

Both watchdog devices are supported in the 32-bit AMD and Intel architecture and AMD64 and Intel 64 architectures for guest operating systems Red Hat Enterprise Linux 6.2 and above.

### 3.3.2. Paravirtualized Devices

Paravirtualization provides a fast and efficient means of communication for guests to use devices on the host machine. KVM provides paravirtualized devices to virtual machines using the virtio API as a layer between the hypervisor and guest.

Some paravirtualized devices decrease I/O latency and increase I/O throughput to near bare-metal levels, while other paravirtualized devices add functionality to virtual machines that is not otherwise available. It is recommended to use paravirtualized devices instead of emulated devices for virtual machines running I/O intensive applications.

All virtio devices have two parts: the host device and the guest driver. Paravirtualized device drivers make it possible for the guest operating system access to physical devices on the host system.

The paravirtualized device drivers must be installed on the guest operating system. By default, the paravirtualized device drivers are included in Red Hat Enterprise Linux 4.7 and later, Red Hat Enterprise Linux 5.4 and later, and Red Hat Enterprise Linux 6.0 and later. The paravirtualized device drivers must be manually installed on Windows guests.



#### Note

For more information on using the paravirtualized devices and drivers, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

#### The paravirtualized network device (virtio-net)

The paravirtualized network device is a virtual network device that provides network access to virtual machines with increased I/O performance and lower latency.

#### The paravirtualized block device (virtio-blk)

The paravirtualized block device is a high-performance virtual storage device with that provides storage to virtual machines with increased I/O performance and lower latency. The paravirtualized block device is supported by the hypervisor and is attached to the virtual machine (except for floppy disk drives, which must be emulated).

#### The paravirtualized controller device (virtio-scsi)

The paravirtualized SCSI controller device is a new feature in Red Hat Enterprise Linux 6.4 that provides a more flexible and scalable alternative to virtio-blk. A **virtio-scsi** guest is capable of inheriting the feature set of the target device, and can handle hundreds of devices compared to virtio-blk, which can only handle 28 devices.

In Red Hat Enterprise Linux 6.4 and above, **virtio-scsi** is fully supported for the following guest operating systems:

- Red Hat Enterprise Linux 6.4 and above
- Windows Server 2008
- Windows 7

- Windows Server 2012
- Windows 8 (32/64 bit)

### The paravirtualized clock

Guests using the Time Stamp Counter (TSC) as a clock source may suffer timing issues. KVM works around hosts that do not have a constant Time Stamp Counter by providing guests with a paravirtualized clock. Additionally, the paravirtualized clock assists with time adjustments needed after a guest runs the sleep (S3) or suspend to RAM operations.

### The paravirtualized serial device (virtio-serial)

The paravirtualized serial device is a bytestream-oriented, character stream device, and provides a simple communication interface between the host's user space and the guest's user space.

### The balloon device (virtio-balloon)

The balloon device can designate part of a virtual machine's RAM as not being used (a process known as *inflating* the balloon), so that the memory can be freed for the host (or for other virtual machines on that host) to use. When the virtual machine needs the memory again, the balloon can be *deflated* and the host can distribute the RAM back to the virtual machine.

### The paravirtualized graphics card (QXL)

The paravirtualized graphics card works with the QXL driver to provide an efficient way to display a virtual machine's graphics from a remote host. The QXL driver is required to use SPICE.

## 3.3.3. Physical Host Devices

Certain hardware platforms allow virtual machines to directly access various hardware devices and components. This process in virtualization is known as *device assignment*, or also as *passthrough*.

### PCI device assignment

The KVM hypervisor supports attaching PCI devices on the host system to virtual machines. PCI device assignment provides guests with exclusive access to PCI devices for a range of tasks. It enables PCI devices to appear and behave as if they were physically attached to the guest virtual machine.

Device assignment is supported on PCI Express devices, with the exception of graphics cards. Parallel PCI devices may be supported as assigned devices, but they have severe limitations due to security and system configuration conflicts.



#### Note

For more information on device assignment, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

### USB passthrough

The KVM hypervisor supports attaching USB devices on the host system to virtual machines. USB device assignment makes it possible for guests to have exclusive access to USB devices for a range of tasks. It also enables USB devices to appear and behave as if they were physically attached to the virtual machine.



### Note

For more information on USB passthrough, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## SR-IOV

SR-IOV (Single Root I/O Virtualization) is a PCI Express standard that extends a single physical PCI function to share its PCI resources as separate, virtual functions (VFs). Each function is capable of being used by a different virtual machine through PCI device assignment.

An SR-IOV-capable PCI-e device, provides a Single Root Function (for example, a single Ethernet port) and presents multiple, separate virtual devices as unique PCI device functions. Each virtual device may have its own unique PCI configuration space, memory-mapped registers, and individual MSI-based interrupts.



### Note

For more information on SR-IOV, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

## NPIV

N\_Port ID Virtualization (NPIV) is a functionality available with some Fibre Channel devices. NPIV shares a single physical N\_Port as multiple N\_Port IDs. NPIV provides similar functionality for Fibre Channel Host Bus Adapters (HBAs) that SR-IOV provides for PCIe interfaces. With NPIV, virtual machines can be provided with a virtual Fibre Channel initiator to Storage Area Networks (SANs).

NPIV can provide high density virtualized environments with enterprise-level storage solutions.



### Note

For more information on NPIV, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## 3.4. Guest CPU Models

*CPU models* define which host CPU features are exposed to the guest operating system. **KVM** and **libvirt** contain definitions for several current processor models, allowing users to enable CPU features that are available only in newer CPU models. The set of CPU features that can be exposed to guests depends on support in the host CPU, the kernel, and **KVM** code.

To ensure safe migration of virtual machines between hosts with different sets of CPU features, **KVM** does not expose all features of the host CPU to guest operating systems by default. Instead, CPU features are based on the selected CPU model. If a virtual machine has a given CPU feature enabled, it cannot be migrated to a host that does not support exposing that feature to guests.



## Note

For more information on guest CPU models, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

## 3.5. Storage

Storage for virtual machines is abstracted from the physical storage used by the virtual machine. It is attached to the virtual machine using the paravirtualized or emulated block device drivers.

### 3.5.1. Storage Pools

A *storage pool* is a file, directory, or storage device managed by **libvirt** for the purpose of providing storage to virtual machines. Storage pools are divided into storage *volumes* that store virtual machine images or are attached to virtual machines as additional storage. Multiple guests can share the same storage pool, allowing for better allocation of storage resources. Refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#) for more information.

#### Local storage pools

Local storage pools are attached directly to the host server. They include local directories, directly attached disks, physical partitions, and Logical Volume Management (LVM) volume groups on local devices. Local storage pools are useful for development, testing and small deployments that do not require migration or large numbers of virtual machines. Local storage pools may not be suitable for many production environments, because they do not support live migration.

#### Networked (shared) storage pools

Networked storage pools include storage devices shared over a network using standard protocols. Networked storage is required when migrating virtual machines between hosts with **virt-manager**, but is optional when migrating with **virsh**. Networked storage pools are managed by **libvirt**.

### 3.5.2. Storage Volumes

Storage pools are divided into storage volumes. Storage volumes are an abstraction of physical partitions, LVM logical volumes, file-based disk images and other storage types handled by **libvirt**. Storage volumes are presented to virtual machines as local storage devices regardless of the underlying hardware.

### 3.5.3. Emulated Storage Devices

The host emulates a range of storage devices that can be presented to virtual machines. Each type of storage device is appropriate for specific use cases. Choice between different types of storage devices allows for maximum flexibility and compatibility with guest operating systems.

#### **virtio-scsi**

**virtio-scsi** is the recommended paravirtualized storage device for guests using large numbers of disks, or advanced storage features such as TRIM.

#### **virtio-blk**

**virtio-blk** is a paravirtualized storage device suitable for exposing image files to guests. **virtio-blk** can provide the best disk I/O performance for virtual machines, but has fewer features than **virtio-scsi**.

## IDE

IDE is recommended for legacy guests that do not support virtio drivers. IDE performance is lower than virtio-scsi or virtio-blk, but it is widely compatible with different systems.

## CD-ROM

ATAPI CD-ROMs and virtio-scsi CD-ROMs are available for presenting ISO files or the host CD-ROM drive to guests. virtio-scsi CD-ROMs can be used with guests that have the virtio-scsi driver installed. ATAPI CD-ROMs offer wider compatibility but lower performance.

## USB mass storage devices and floppy disks

Emulated USB mass storage devices and floppy disks are available when removable media is required. USB mass storage devices are preferable to floppy disks due to their larger capacity.

### 3.5.4. Host Storage

Disk images can be stored on a range of local and remote storage technologies connected to the host.

#### Image files

Image files can only be stored on a host file system. The image files can be stored on a local file system, such as ext4 or xfs, or a network file system, such as NFS.

Tools such as **libguestfs** can manage, back up, and monitor files. Disk image formats on KVM include:

##### **raw**

Raw image files contain the contents of the disk with no additional metadata.

Raw files can either be pre-allocated or sparse, if the host file system allows it. Sparse files allocate host disk space on demand, and are therefore a form of thin provisioning. Pre-allocated files are fully provisioned but have higher performance than sparse files.

Raw files are desirable when disk I/O performance is critical and transferring the image file over a network is rarely necessary.

##### **qcow2**

qcow2 image files offer a number of advanced disk image features, including backing files, snapshots, compression, and encryption. They can be used to instantiate virtual machines from template images.

qcow2 files are typically more efficient to transfer over a network, because only sectors written by the virtual machine are allocated in the image.

#### LVM volumes

Logical volumes (LVs) can be used for disk images and managed using the system's LVM tools. LVM offers higher performance than file systems because of its simpler block storage model.

LVM thin provisioning offers snapshots and efficient space usage for LVM volumes, and can be used as an alternative to migrating to qcow2.

#### Host devices

Host devices such as physical CD-ROMs, raw disks, and logical unit numbers (LUNs) can be presented to the guest. This enables a guest to use storage area network (SAN) or iSCSI LUNs, as well as local CD-ROM media, with good performance.

Host devices can be used when storage management is done on a SAN instead of on hosts.

### Distributed storage systems

Gluster volumes can be used as disk images. This enables high-performance clustered storage over the network.

Red Hat Enterprise Linux 6.5 and above includes native support for creating virtual machines with GlusterFS. This enables a KVM host to boot virtual machine images from GlusterFS volumes, and to use images from a GlusterFS volume as data disks for virtual machines. When compared to GlusterFS FUSE, the native support in KVM delivers higher performance.



#### Note

For more information on storage and virtualization, refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## 3.6. Virtual Networking

A virtual guest's connection to any network uses the software network components of the physical host. These software components can be rearranged and reconfigured by using **libvirt**'s virtual network configuration. The host therefore acts as a [virtual network switch](#), which can be configured in a number of different ways to fit the guest's networking needs.

By default, all guests on a single host are connected to the same libvirt virtual network, named **default**. Guests on this network can make the following connections:

#### With each other and with the virtualization host

Both inbound and outbound traffic is possible, but is affected by the firewalls in the guest operating system's network stack and by [libvirt network filtering](#) rules attached to the guest interface.

#### With other hosts on the network beyond the virtualization host

Only outbound traffic is possible, and is affected by [Network Address Translation \(NAT\)](#) rules, as well as the host system's firewall.

However, if needed, guest interfaces can instead be set to one of the following modes:

#### [Isolated mode](#)

The guests are connected to a network that does not allow any traffic beyond the virtualization host.

#### [Routed mode](#)

The guests are connected to a network that routes traffic between the guest and external hosts without performing any NAT. This enables incoming connections but requires extra routing-table entries for systems on the external network.

#### [Bridged mode](#)

The guests are connected to a bridge device that is also connected directly to a physical ethernet device connected to the local ethernet. This makes the guest directly visible on the physical network, and thus enables incoming connections, but does not require any extra routing-table entries.

For basic outbound-only network access from virtual machines, no additional network setup is usually needed, as the **default** network is installed along with the *libvirt* package, and automatically started when the **libvirtd** service is started. If more advanced functionality is needed, additional networks can be created and configured using either [virsh](#) or [virt-manager](#), and the [guest XML configuration file](#) can be edited to use one of these new networks.



### Note

For information on advanced virtual network settings, see the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

From the point of view of the guest operating system, a virtual network connection is the same as a normal physical network connection. For further information on configuring networks in Red Hat Enterprise Linux 6 guests, see the [Red Hat Enterprise Linux 6 Deployment Guide](#).

## Chapter 4. Quick Start Tutorial on Virtualization in Red Hat Enterprise Linux 6

This chapter provides instructions for installing essential virtualization tools and creating a virtual machine in Red Hat Enterprise Linux 6.



### Note

The tutorial uses Virtual Machine Manager to quickly create a virtual machine for trying out KVM virtualization. To set up virtual machines with the capabilities necessary for a production environment, refer to the [Red Hat Enterprise Linux 6 Host Configuration and Guest Installation Guide](#).

### 4.1. Tutorial overview

- » [Ensure that your system meets the basic requirements](#)
- » [Ensure that the required packages are installed](#)
- » [Create a virtual machine using Virtual Machine Manager](#)

### 4.2. Basic Requirements and Setup

To set up a KVM virtual machine on Red Hat Enterprise Linux 6, your system must meet the following criteria:

#### Architecture

Virtualization with the KVM hypervisor is currently only supported on Intel 64 and AMD64 systems.

#### Disk space and RAM

Minimum:

- » 6 GB free disk space
- » 2 GB RAM

#### Customer Portal subscription

To install virtualization packages, your host machine must be registered and subscribed to the Red Hat Customer Portal. To register run the **subscription-manager register** command and follow the prompts. Alternatively, run the Red Hat Subscription Manager application from **Applications** → **System Tools** on the desktop to register.

If you do not have a valid Red Hat subscription, visit the [Red Hat online store](#) to obtain one. For more information on registering and subscribing a system to the Red Hat Customer Portal, see <https://access.redhat.com/solutions/253273>.

#### Required packages

Before you can use virtualization, a basic set of virtualization packages must be installed on your computer.

#### Procedure 4.1. Installing the virtualization packages with yum



To use virtualization on Red Hat Enterprise Linux the *libvirt*, *qemu-kvm* and *qemu-img* packages must be installed. These packages provide the user-level KVM emulator and disk image manager on the host system.

1. Install the *qemu-kvm*, *qemu-img*, *libvirt*, and *virt-manager* packages with the following command:

```
# yum install qemu-kvm qemu-img libvirt virt-manager
```

2. Download a Red Hat Enterprise Linux 7 Workstation binary DVD ISO image from the [Red Hat Customer Portal](#). This image will be used to install the guest virtual machine's operating system.



### Note

If you encounter any problems during the installation process, see the [Troubleshooting section](#) of the Red Hat Enterprise Linux 6 *Virtualization Host Configuration and Guest Installation Guide*.

## 4.3. Creating a Virtual Machine with Virtual Machine Manager

The **Virtual Machine Manager**, also known as **virt-manager**, is a graphical tool for quick deployment of virtual machines in Red Hat Enterprise Linux. In this tutorial, you will become familiar with its basic functions and will be able to use **Virtual Machine Manager** to create a virtual machine.

### 4.3.1. Introduction to Virtual Machine Manager

To open the **Virtual Machine Manager**, click **Applications** → **System Tools** → **Virtual Machine Manager**; or open the terminal and run the **virt-manager** command.

The following image shows the **Virtual Machine Manager** interface. This interface enables you to control all of your virtual machines from one central location.

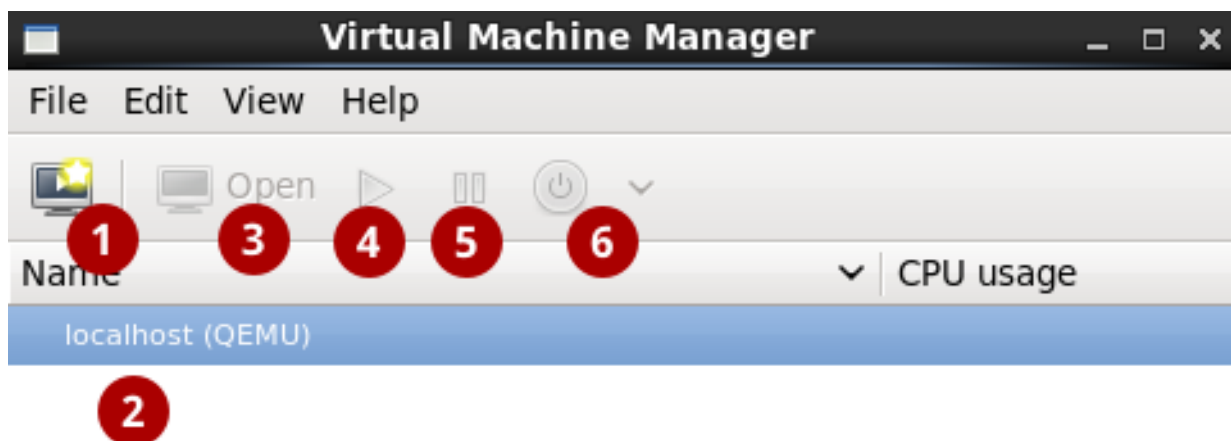


Figure 4.1. The Virtual Machine Manager interface

Commonly used interface elements include:

- **1 Create new virtual machine:** Click here to create a new virtual machine.
- **2 Virtual machines:** A list of configured connections and all guest virtual machines associated with them. When a virtual machine is created, it will be listed here. When a guest is running, an animated graph shows the guest's CPU usage in the **CPU usage** column.

After selecting a virtual machine from this list, use the following buttons to control the selected virtual machine's state:

- **3 Open:** Opens the guest virtual machine console and details in a new window.
- **4 Run:** Turns on the virtual machine.
- **5 Pause:** Pauses the virtual machine.
- **6 Shut down:** Shuts down the virtual machine. Clicking on the arrow displays a drop-down menu with several options for turning off the virtual machine, including Reboot, Shut Down, Force Reset, Force Off, and Save.

Right-clicking on a virtual machine shows a menu with more functions, including:

- **Clone:** Clones the virtual machine.
- **Migrate:** Migrates the virtual machine to another host.
- **Delete:** Deletes the virtual machine.

### 4.3.2. Creating a Virtual Machine with Virtual Machine Manager

Follow these steps to create a Red Hat Enterprise Linux 7 virtual machine on **Virtual Machine Manager**.

#### Procedure 4.2. Creating a guest virtual machine with Virtual Machine Manager

##### 1. Open Virtual Machine Manager

Click **Applications** → **System Tools** → **Virtual Machine Manager**

or

Open the terminal and use the **virt-manager**.

##### 2. Create a new virtual machine



Click  to open the **New VM** wizard.

### 3. Specify name and installation method

In Step 1, type in a virtual machine name and choose an installation type to install the guest virtual machine's operating system.

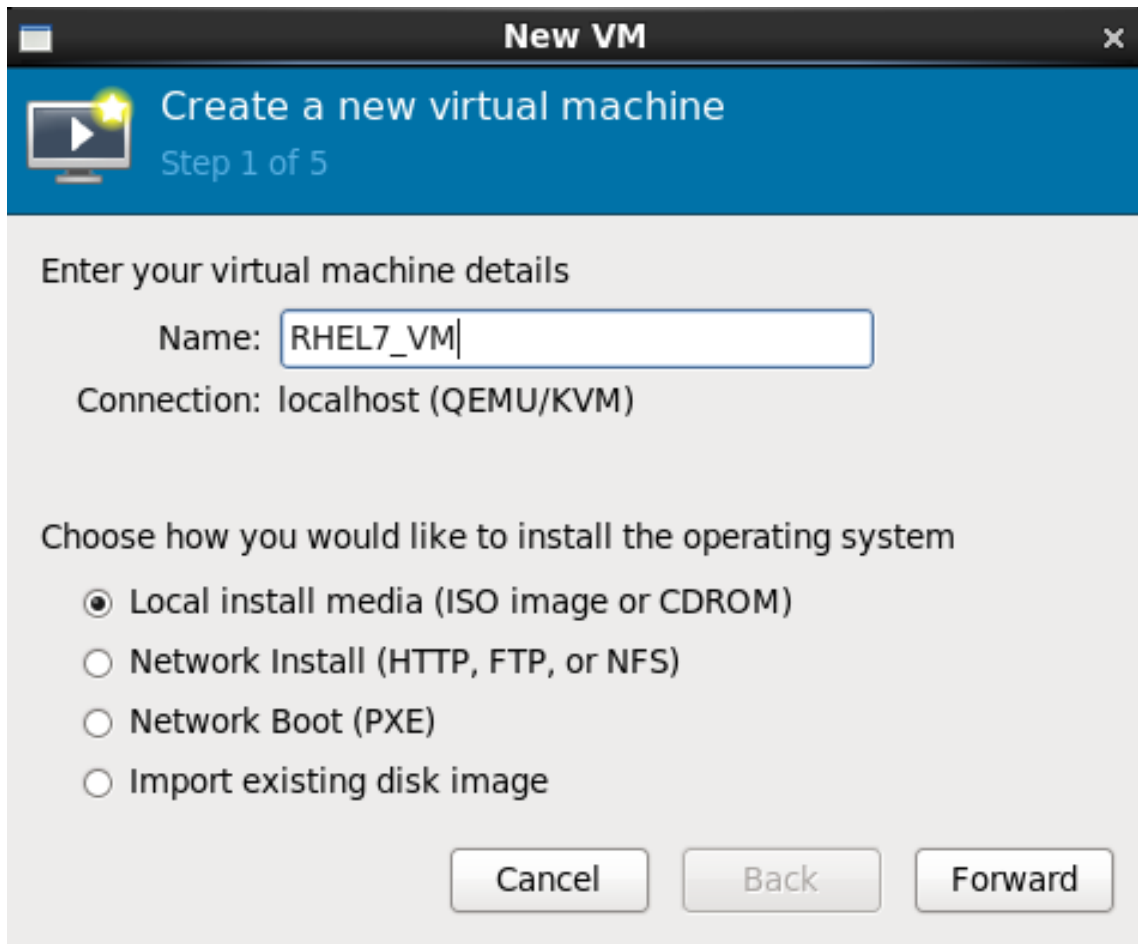


Figure 4.2. Name virtual machine and select installation method

For this tutorial, select **Local install media (ISO image)**. This installation method uses an image of an installation disk (in this case an `.iso` file). Click **Forward** to continue to the next step.

### 4. Locate installation media

- a. Select the **Use ISO Image** option.
- b. Click **Browse** → **Browse Local** buttons.
- c. Locate the ISO downloaded in [Procedure 4.1, “Installing the virtualization packages with yum”](#) on your machine.
- d. Select the ISO file and click **Open**.
- e. Ensure that Virtual Machine Manager correctly detected the OS type. If not, uncheck **Automatically detect operating system based on install media** and select **Linux** from the **OS type** drop-down and **Red Hat Enterprise Linux 7** from the **Version** drop-down.

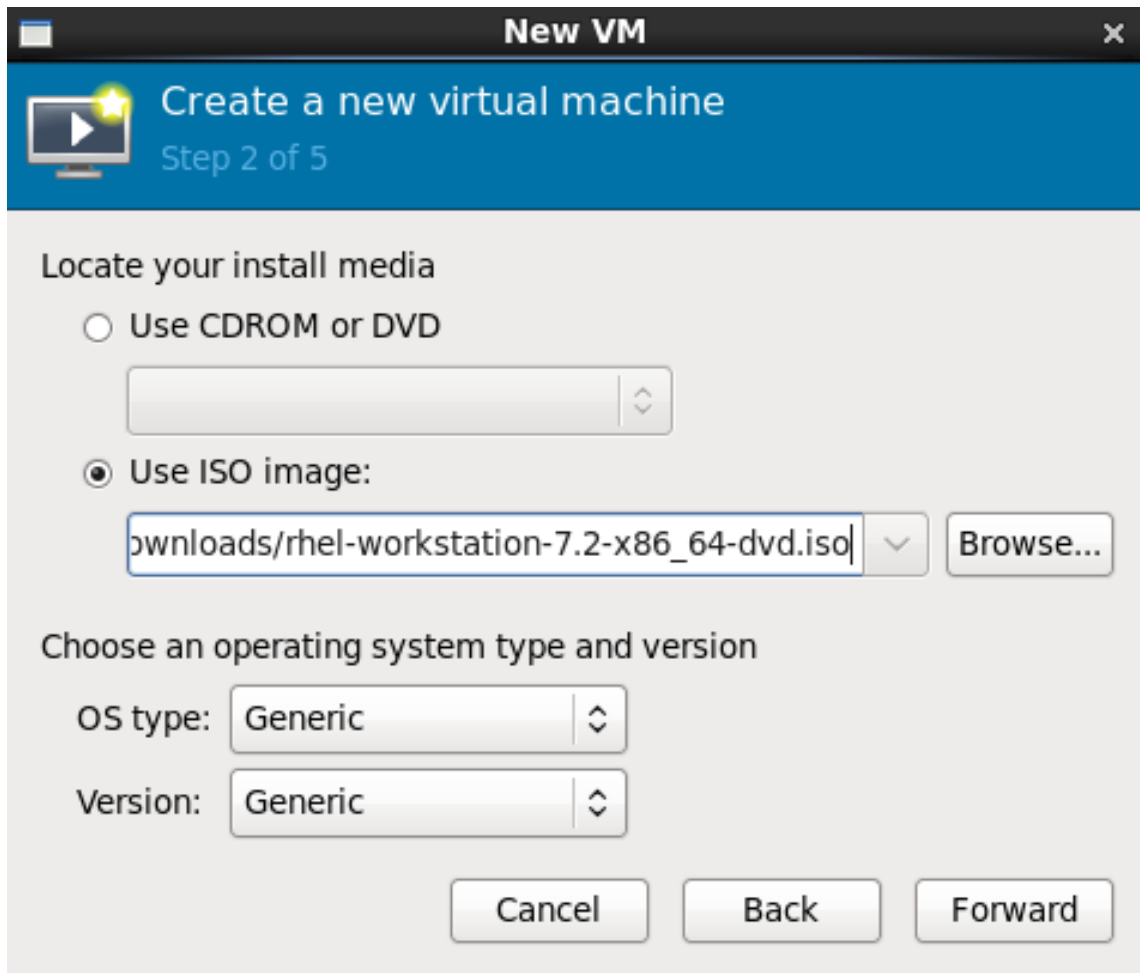


Figure 4.3. Local ISO image installation

## 5. Configure memory and CPU

You can use step 3 of the wizard to configure the amount of memory and the number of CPUs to allocate to the virtual machine. The wizard shows the number of CPUs and amount of memory available to allocate.

For this tutorial, leave the default settings and click **Forward**.

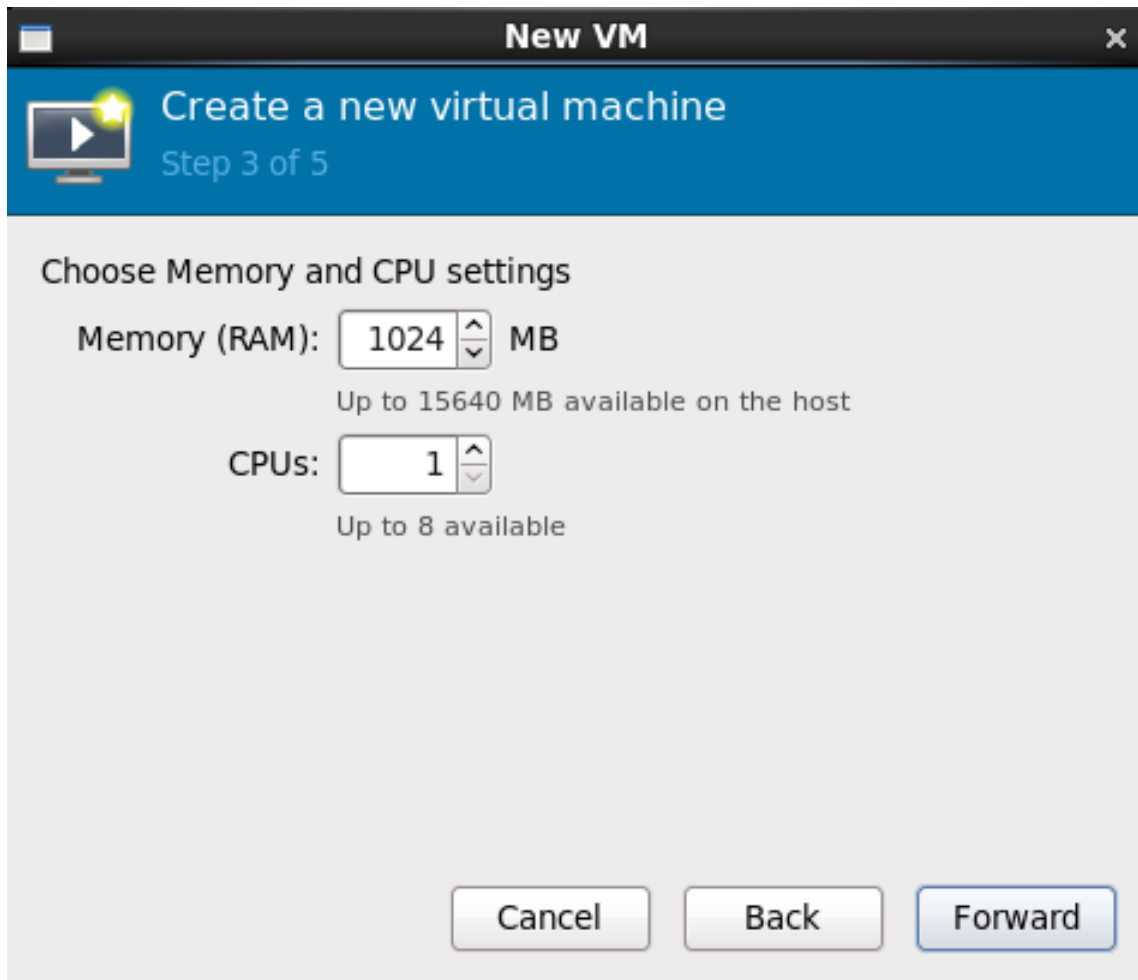


Figure 4.4. Configuring CPU and Memory

## 6. Configure storage

Using step 4 of the wizard, you can assign storage to the guest virtual machine. The wizard shows options for storage, including where to store the virtual machine on the host machine. For this tutorial, leave the default settings and click **Forward**.

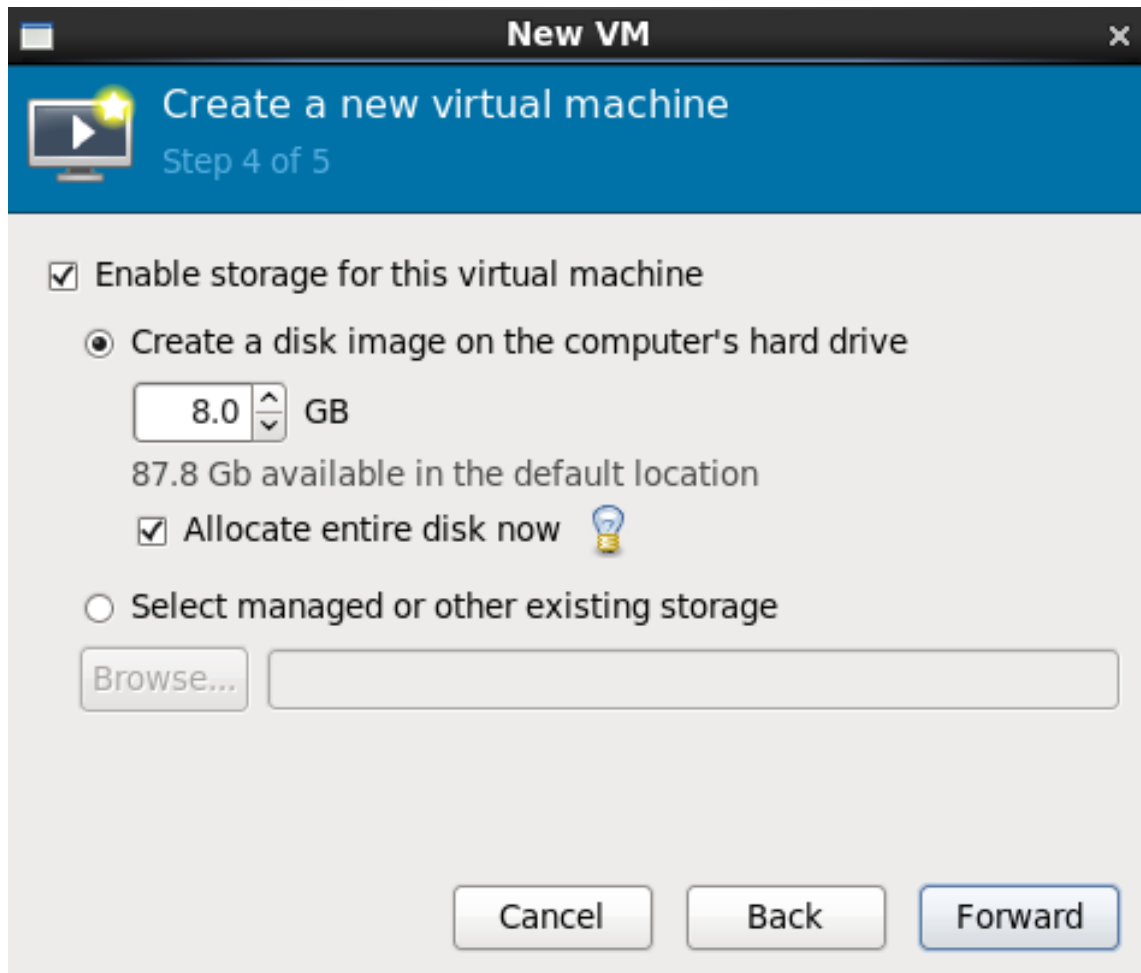


Figure 4.5. Configuring storage

## 7. Review the configuration

Using step 5 of the wizard, you can configure the virtualization type, and guest architecture, and networking settings. For this tutorial, verify the settings, and click **Finish**. **Virtual Machine Manager** will create a virtual machine with the specified hardware settings.

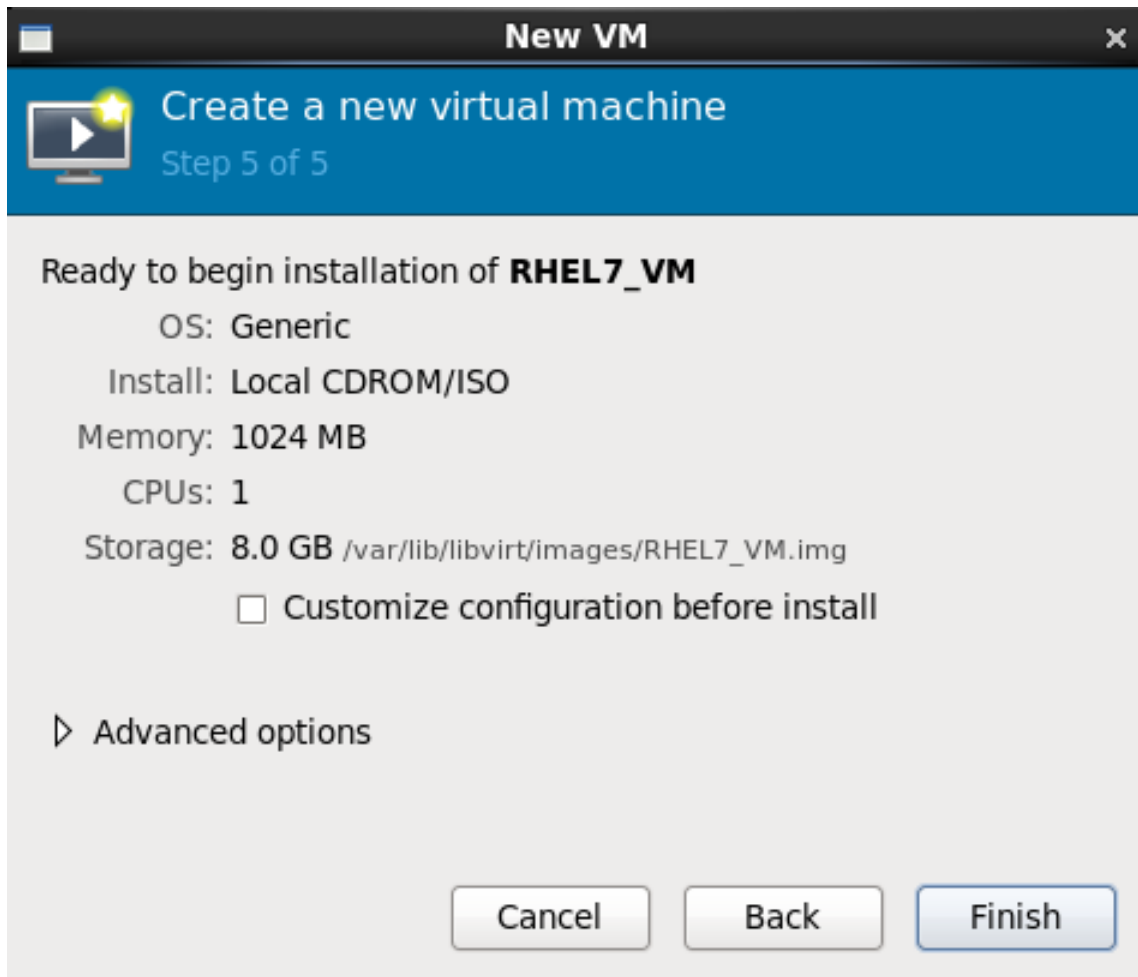


Figure 4.6. Verifying the configuration

After **Virtual Machine Manager** creates your Red Hat Enterprise Linux 7 virtual machine, the virtual machine's window will open, and the installation of the selected operating system will begin in it. Follow the instructions in the Red Hat Enterprise Linux 7 installer to complete the installation of the virtual machine's operating system.



### Note

For help with Red Hat Enterprise Linux 7 installation, refer to the [Red Hat Enterprise Linux 7 Installation Guide](#).

### 4.3.3. Exploring the Guest Virtual Machine

You can view a virtual machine's console by selecting a virtual machine in the **Virtual Machine Manager** window and clicking **Open**. You can operate your Red Hat Enterprise Linux 7 virtual machine from the console in the same way as a physical system.




Figure 4.7. The guest virtual machine console

- ✦ **1 Show the graphical console:** Shows the virtual machine's display. The virtual machine can be operated from the console the same as a physical machine.
- ✦ **2 Show virtual hardware details:** Shows details about the virtual hardware that the guest is using. These include an overview of basic system details, performance, processor, memory, and boot settings, and details of the system's virtual devices.
- ✦ These buttons control the virtual machine's state:
  - **3 Run:** Turns on the virtual machine.
  - **4 Pause:** Pauses the virtual machine.
  - **5 Shut down:** Shuts down the virtual machine. Clicking on the arrow displays a drop-down menu with several options for turning off the virtual machine, including Reboot, Shut Down, Force Reset, Force Off, and Save.



- ✦ **6 Full screen:** Switches the virtual machine to full screen view.
- ✦ **7 Send Key:** Sends key combinations such as Ctrl+Alt+Backspace, Ctrl+Alt+Delete, Ctrl+Alt+F1, PrintScreen, and more to the virtual machine.



### Note

For more information about using the **Virtual Machine Manager** to create and run virtual machines, refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#) and the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#).

## Chapter 5. Virtualization Tools

This chapter provides an introduction to the basic tools available to interact with virtualization on Red Hat Enterprise Linux 6.

### 5.1. virsh

*virsh* is a command-line interface (CLI) tool for managing the hypervisor and guest virtual machines. **virsh** is built on the **libvirt** management API. . The **virsh** command can be used in read-only mode by unprivileged users, or for full administration by users with root access. In addition, **virsh** is the main management interface for **virsh** guest domains and can be used to create, pause, and shut down domains, as well as list current domains. This tool is installed as part of the *libvirt-client* package.



#### Note

Refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#) for more information about managing virtual machines with **virsh**.

### 5.2. virt-manager

*virt-manager* is a lightweight graphical tool for managing virtual machines, more user-friendly and less powerful than **virsh**. It provides the ability to control the life cycle of existing machines, provision new machines, manage virtual networks, access the graphical console of virtual machines, and view performance statistics. For a demonstration of using *virt-manager*, see [Section 4.3, “Creating a Virtual Machine with Virtual Machine Manager”](#). This tool is provided by the *virt-manager* package.



#### Note

Refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#) for more information about managing virtual machines with **virt-manager**.

### 5.3. virt-install

*virt-install* is a command-line tool for provisioning new virtual machines. It supports both text-based and graphical installations, using serial console, SDL, SPICE, or VNC client-server pair graphics. Installation media can be local, or exist remotely on an NFS, HTTP, or FTP server. The tool can also be configured to run unattended and kickstart the guest when installation is complete, allowing for easy automation of installation. This tool is installed as part of the *python-virtinst* package.



#### Note

Refer to the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#) for more information about **virt-install**.

### 5.4. guestfish

## 5.4. **guestfish**

*guestfish* is a shell and command-line tool for examining and modifying virtual machine disk images. This tool uses *libguestfs* and exposes all functionality provided by the **guestfs** API.



### Warning

Using **guestfish** on running virtual machines may cause disk-image corruption. Use the **guestfish** command with the **--ro** (read-only) option if the disk image is being used by a running virtual machine.



### Note

Refer to the [Red Hat Enterprise Linux 6 Virtualization Administration Guide](#) for more information about **guestfish**.

## 5.5. Other Useful Tools

The following tools are used to access a guest virtual machine's disk from the host. The guest's disk is usually accessed directly using the **disk-image** file located on the host. The commands that follow are part of the **libvirt** domain and are used to gain access to the guest's disk image.

### **guestmount**

A command-line tool used to mount virtual machine file systems and disk images on the host machine.



### Warning

Using **guestmount** in **--r/w** (read/write) mode to access a disk that is currently being used by a guest may cause the disk to become corrupted. Do not use **guestmount** in **--r/w** (read/write) mode on live virtual machines. Use the **guestmount** command with the **--ro** (read-only) option if the disk image is being used.

### **virt-cat**

A command-line tool that can be used to quickly view the contents of one or more files in a specified virtual machine's disk or disk image. This tool is installed as part of the *libguestfs-tools-c* package.

### **virt-df**

A command-line tool used to show the actual physical disk usage of virtual machines, similar to the command-line tool **df**. Note that this tool does not work across remote connections. It is installed as part of the *libguestfs-tools-c* package.

### **virt-edit**

A command-line tool used to edit files that exist on a specified virtual machine. This tool is installed as part of the *libguestfs-tools-c* package.



## Warning

Using **virt-edit** on live virtual machines may cause disk corruption in the virtual machine. Although the **virt-edit** command will try to prevent users from editing files on live virtual machines, it is not guaranteed to catch all instances. Do not use **virt-edit** on a live virtual machine.

## virt-filesystems

A command-line tool used to discover file systems, partitions, logical volumes and their sizes in a disk image or virtual machine. One common use is in shell scripts, to iterate over all file systems in a disk image. This tool is installed as part of the *libguestfs-tools-c* package.



## Note

This tool replaces **virt-list-filesystems** and **virt-list-partitions**.

## virt-inspector

A command-line tool that can examine a virtual machine or disk image to determine the version of its operating system and other information. It can also produce XML output, which can be piped into other programs. Note that **virt-inspector** can only inspect one domain at a time. This tool is installed as part of the *libguestfs-tools-c* package.

## virt-inspector2

An alternative tool to **virt-inspector**, written in C. This tool is installed as part of the *libguestfs-tools-c* package.

## virt-ls

A command-line tool that lists files and directories inside a virtual machine. This tool is installed as part of the *libguestfs-tools-c* package.

## virt-make-fs

A command-line tool for creating a file system based on a tar archive or files in a directory. It is similar to tools like **mkisofs** and **mksquashfs**, but it can create common file system types such as ext2, ext3 and NTFS, and the size of the file system created can be equal to or greater than the size of the files it is based on. This tool is provided as part of the *libguestfs-tools* package.

## virt-p2v

A graphical tool to convert physical machines into virtual machines. This tool is installed as part of the *virt-v2v* package in Red Hat Enterprise Linux 6.2 and later.

## virt-rescue

A command-line tool that provides a rescue shell and some simple recovery tools for unbootable virtual machines and disk images. It can be run on any virtual machine known to **libvirt**, or directly on disk images. This tool is installed as part of the *libguestfs-tools-c* package.



### Warning

Using **virt-rescue** on running virtual machines may cause disk corruption in the virtual machine. **virt-rescue** attempts to prevent its own use on running virtual machines, but cannot catch all cases. Using the command with the **--ro** (read-only) option will not cause disk corruption, but may give strange or inconsistent results.

Avoid using **virt-rescue** on a running virtual machine.

### virt-resize

A command-line tool to resize virtual machine disks, and resize or delete any partitions on a virtual machine disk. It works by copying the guest image and leaving the original disk image untouched. This tool is installed as part of the *libguestfs-tools-c* package.



### Important

Using **virt-resize** on running virtual machines can give inconsistent results. It is best to shut down virtual machines before attempting to resize them.

### virt-sysprep

A command-line tool to reset, customize, or unconfigure virtual machines to prepare a template for creating clones. This tool is installed as part of the *libguestfs-tools-c* package.



### Important

Virtual machines must be shut down before running **virt-sysprep**. To preserve a virtual machine's existing contents, snapshot, copy or clone the disk before running **virt-sysprep**.

### virt-tar

A command-line archive tool for downloading and uploading parts of a virtual machine's file system. This tool is commonly used for making backups, uploading data, reviewing guest activity, and fixing or customizing guests. This tool is installed as part of the *libguestfs-tools* package.



### Warning

Using **virt-tar** with the **-u** (upload) option on running virtual machines can cause disk corruption in the virtual machine. **virt-tar** attempts to prevent its own use on running virtual machines, but cannot catch all cases.

Using **virt-tar** with the **-x** (extract) option on running virtual machines will not cause disk corruption, but may give strange or inconsistent results. It is best to shut down virtual machines before attempting to extract files from them.

## **virt-top**

A command-line utility similar to **top**, which shows statistics related to virtualized domains. This tool is provided by the *virt-top* package.

## **virt-v2v**

A graphical tool to convert virtual machines from Xen and VMware hypervisors to run on KVM. This tool is provided by the *virt-v2v* package.

## **virt-viewer**

A lightweight utility for displaying the graphical console of a virtual machine using the VNC and SPICE protocols. This tool is provided by the *virt-viewer* package.

## **virt-what**

A shell script that detects whether a program is running in a virtual machine. This tool is provided by *virt-what* package.

## **virt-who**

The *virt-who* package is a Red Hat Enterprise Linux host agent that queries **libvirt** for guest UUIDs. It then passes that data to the local entitlement server for the purposes of issuing certificates. This tool is provided by the *virt-who* package.



### Note

For further information on using **virt-who**, see the [Red Hat Enterprise Linux 6 Virtualization Host Configuration and Guest Installation Guide](#).

## **virt-win-reg**

A command-line tool to export and merge Windows Registry entries from a Windows virtual machine, and perform simple Registry operations. This tool is installed as part of the *libguestfs-tools* package.



### Warning

Using **virt-win-reg** on running virtual machines will cause irreversible disk corruption in the virtual machine. **virt-win-reg** attempts to prevent its own use on running virtual machines, but cannot catch all cases.



### Warning

Modifying the Windows Registry is an inherently risky operation, as the format is deliberately obscure and undocumented. Changes to the registry can leave the system unbootable, so ensure you have a reliable backup before you use the **--merge** option.

## **virt-xml-validate**

A command-line tool to validate **libvirt** XML files for compliance with the published schema. This tool is installed as part of the *libvirt-client* package.

## Appendix A. Revision History

<b>Revision 1.1-10</b>	<b>Wed Mar 08 2017</b>	<b>Jiri Herrmann</b>
Updates for the 6.9 GA release		
<b>Revision 1.1-8</b>	<b>Mon December 20 2016</b>	<b>Jiri Herrmann</b>
Updates for the 6.9 beta release		
<b>Revision 1.1-7</b>	<b>Mon May 02 2016</b>	<b>Jiri Herrmann</b>
Updates for the 6.8 GA release		
<b>Revision 1.1-6</b>	<b>Tue Mar 01 2016</b>	<b>Jiri Herrmann</b>
Prepared the book for 6.8 beta publication		
<b>Revision 1.1-5</b>	<b>Thu Oct 08 2015</b>	<b>Jiri Herrmann</b>
Cleaned up the Revision History		
<b>Revision 1.1-4</b>	<b>Mon Jul 13 2015</b>	<b>Jiri Herrmann</b>
Version for 6.7 GA release		
<b>Revision 1.1-03</b>	<b>Thurs Apr 16 2015</b>	<b>Dayle Parker</b>
Preparing document for 6.7 Beta publication.		