



Red Hat Enterprise Linux 7 Virtualization Getting Started Guide

An introduction to virtualization concepts

Dayle Parker
Jacquelynn East

Laura Novich
Scott Radvan

Jiri Herrmann

Red Hat Enterprise Linux 7 Virtualization Getting Started Guide

An introduction to virtualization concepts

Dayle Parker
Red Hat Customer Content Services
dayleparker@redhat.com

Laura Novich
Red Hat Customer Content Services
lnovich@redhat.com

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

Jacquelynn East
Red Hat Customer Content Services
jeast@redhat.com

Scott Radvan
Red Hat Customer Content Services

Legal Notice

Copyright © 2011-2015 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](#). 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, MetaMatrix, 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

Chapter 1. What is virtualization and migration?	2
1.1. What is virtualization?	2
1.2. What is migration?	2
Chapter 2. Advantages and misconceptions of virtualization	4
2.1. Virtualization costs	4
2.2. Virtualization learning curve	4
2.3. Performance	4
2.4. Disaster recovery	5
2.5. Security	5
2.6. Virtualization for servers and individuals	6
Chapter 3. Introduction to Red Hat virtualization products	7
3.1. KVM and virtualization in Red Hat Enterprise Linux	7
3.2. libvirt and libvirt tools	9
3.3. Virtualized hardware devices	10
3.4. Storage	16
Chapter 4. Virtualization Tools	19
4.1. virsh	19
4.2. virt-manager	19
4.3. virt-install	19
4.4. guestfish	20
4.5. GNOME Boxes	20
4.6. Other useful tools	20
Chapter 5. Quick Start Guide to Virtualization in Red Hat Enterprise Linux 7	25
5.1. Basic requirements and setup	25
5.2. Deploying a virtual machine with Virtual Machine Manager	26
Appendix A. Revision History	36

Chapter 1. What is virtualization and migration?

This chapter discusses terms related to virtualization and migration.

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. Most existing implementations of virtualization use a *hypervisor* - a software layer or subsystem that controls hardware and allows multiple operating systems, called *guests* - to run on the same physical system, called a *host*. There are several methods for virtualizing operating systems.

Virtualization methods

Full virtualization

Full virtualization uses the hardware capabilities of the processor to allow guests to use features of the host's physical system. This creates a new virtual system, called a *virtual machine*, which allows guest operating systems to run without modifications. The guest operating system and any applications on the guest virtual machine are unaware of their virtualized environment and run as if they were on a physical machine. The KVM (Kernel-based Virtual Machine) hypervisor in Red Hat Enterprise Linux uses full virtualization, provided by the hardware-assisted virtualization method.

Para-virtualization

Para-virtualization employs a collection of software and data structures that are presented to the virtualized guest, requiring software modifications in the guest to use the para-virtualized environment. Para-virtualization can encompass the entire kernel, as is the case for Xen para-virtualized guests, or drivers that virtualize I/O devices.

Software virtualization (or emulation)

Software virtualization uses slower binary translation and other emulation techniques to run unmodified operating systems. Software virtualization is unsupported by Red Hat Enterprise Linux.



Note

For more information and detailed instructions on guest installation, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

1.2. What is 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 or shuts down 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 Deployment and Administration Guide](#) for details.

It is important to understand that the migration process moves the virtual machine's memory, and the disk volume associated with the virtual machine is also migrated. This process is done using live block migration — information about this can be found in the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

1.2.1. Benefits of migrating virtual machines

Migration is useful for:

Load balancing

When a host machine is overloaded, one or many of its virtual machines could be migrated to other hosts using live migration.

Upgrading or making changes to the host

When the need arises to upgrade, add, or remove hardware devices on one 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 any of the 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 another physical location for lower latency or for other special circumstances.

Chapter 2. Advantages and misconceptions of virtualization

There are many advantages to virtualization and perhaps an equal amount of misconceptions surrounding it. This chapter explores these points.

2.1. Virtualization cost

A common misconception is that virtualization is too expensive to justify the change. Virtualization can be expensive to introduce but often it saves money in the long term. It is important to perform a Return on Investment (ROI) analysis to determine the best use of virtualization in your environment. 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 adequate planning is performed before migrating physical systems to virtualized ones, less time is spent maintaining them. This means less money being spent on parts and labor.

Extended life for installed software

Older versions of software may not run on newer, bare-metal machines directly. However, by running the older software virtually on a larger, faster system, the life of the software may be extended while taking advantage of the performance from the 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. This means the space normally occupied by server hardware can be used for other purposes.

2.2. Virtualization learning curve

A misconception exists that virtualization is difficult to learn. In truth, virtualization is no more difficult or easy to learn than any new process. The skills required for managing and supporting a physical environment are easily transferable to a virtual one. Virtual environments function similarly to their physical counterparts, ensuring the learning curve remains a slight one.

2.3. Performance

On older virtualization versions that supported only a single CPU, 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:

- The industry standard SAP Sales and Distribution (SD) Standard Application Benchmark found Red Hat Enterprise Linux 6.2 and KVM to demonstrate a virtualization efficiency of 85% when comparing 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 details on these virtualization benchmarks, visit:

- Red Hat Knowledgebase, *SAP-SD Benchmark running in a VM – Leadership Performance using RHEL 6 / KVM* at <https://access.redhat.com/knowledge/articles/216943>
- The *Standard Performance Evaluation Corporation (SPEC)* at <http://www.spec.org>
- *Red Hat Achieves New Top Virtualization Performance Benchmark with HP* at <http://investors.redhat.com/releasedetail.cfm?ReleaseID=617594>

2.4. Disaster recovery

Disaster recovery is quicker and easier when the systems are virtualized. On a physical system, if something serious goes wrong, a complete re-install 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 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 from each other, one system's downtime will not affect any others.

2.5. Security

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

2.5.1. Virtualization security features

SELinux

SELinux was developed by the US National Security Agency and others to provide Mandatory Access Control (MAC) for Linux. Under control of SELinux, all processes and files are given what is known as a *type*, and access is limited by fine-grained controls. 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 7.

sVirt

sVirt is a technology included in Red Hat Enterprise Linux 7 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 bugs in the hypervisor that might be used as an attack vector for the host or to another virtual machine.



Note

For more information on security in virtualization, refer to the *Red Hat Enterprise Linux 7 Virtualization Security Guide*.

2.6. Virtualization for servers and individuals

Virtualization is not just for servers; it can be useful for individuals as well. Desktop virtualization offers centralized management, an improved desktop solution, and better disaster recovery. By using connection software, it is possible to connect to a desktop remotely.

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

2.6.1. Virtualization deployment scenarios

These are examples of common deployment scenarios for virtualization, and the tools that can be used to facilitate these scenarios.

Small deployments of up to 3 physical hosts and 10 guests: virt-manager

A tool such as virt-manager can be useful to a small business running several servers that do not have strict uptime requirements or service-level agreements (SLAs). In this environment, a single administrator may be responsible for the entire infrastructure, and maintaining procedural flexibility is important if a component needs to be changed. This environment may contain applications such as web servers, file and print servers, and application servers.

Large deployments or mission-critical applications: Red Hat Enterprise Virtualization (RHEV)

A full virtualization platform such as Red Hat Enterprise Virtualization (RHEV) might suit an enterprise running larger deployments or mission-critical applications. In this environment, the physical infrastructure is large enough to require an IT department and the business requirements demand a defined response to new needs. Some examples of a large deployment suited to Red Hat Enterprise Virtualization may include databases, trading platforms, or messaging systems that must run continuously without any downtime.

Software developers producing management applications: libvirt

Both virt-manager and Red Hat Enterprise Virtualization (RHEV) use libvirt to manage virtual machines. libvirt is a virtualization application programming interface (API) that allows software developers to produce and adapt management applications.

Chapter 3. Introduction to Red Hat virtualization products

This chapter introduces the various virtualization products available in Red Hat Enterprise Linux.

3.1. KVM and virtualization in Red Hat Enterprise Linux

What is KVM?

KVM (Kernel-based Virtual Machine) is a full virtualization solution for Linux on AMD64 and Intel 64 hardware that is built into the standard Red Hat Enterprise Linux 7 kernel and 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.

Overcommitting

The KVM hypervisor supports *overcommitting* of system resources. Overcommitting means allocating more virtualized CPUs or memory than the available resources on the system. Memory overcommitting allows hosts to utilize memory and virtual memory to increase guest densities.



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 7 Virtualization Deployment and Administration Guide*.

Thin provisioning

Thin provisioning allows the allocation of flexible storage and optimizes the available space for every guest virtual machine. It gives the appearance that there is more physical storage on the guest than is actually available. This is not the same as overcommitting as this only pertains to storage and not CPUs or memory allocations.



Important

Like overcommitting, thin provisioning involves possible risks to system stability. For more information on thin provisioning with KVM, and the precautions that should be taken, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

KSM

Kernel Same-page Merging (KSM), used by the KVM hypervisor, allows 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 7 Virtualization Deployment and Administration Guide*.

QEMU Guest Agent

The *QEMU Guest Agent* runs on the guest operating system and allows 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 7 Virtualization Deployment and Administration Guide*.

Disk I/O throttling

When several virtual machines are running simultaneously, they can interfere with system performance by using excessive disk I/O. *Disk I/O throttling* in KVM provides the ability to set a limit on disk I/O requests sent from virtual machines to the host machine. This can prevent a virtual machine from over utilizing shared resources, and impacting the performance of other virtual machines.

**Note**

For instructions on using disk I/O throttling, refer to the *Red Hat Enterprise Linux 7 Virtualization Tuning and Optimization Guide*.

Automatic NUMA balancing

Automatic NUMA balancing improves the performance of applications running on non-uniform memory access (NUMA) hardware systems, without any manual tuning required for Red Hat Enterprise Linux 7 guests. Automatic NUMA balancing moves tasks, which can be threads or processes, closer to the memory they are accessing.

**Note**

For more information on automatic NUMA balancing, refer to the *Red Hat Enterprise Linux 7 Virtualization Tuning and Optimization Guide*.

Virtual CPU hot add

Virtual CPU (vCPU) hot add capability provides the ability to increase processing power as needed on running virtual machines, without downtime. The vCPUs assigned to a virtual machine can be added to a running guest in order to either meet the workload's demands, or to maintain the Service Level Agreement (SLA) associated with the workload.



Note

For more information on virtual CPU hot add, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

KVM guest virtual machine compatibility

Red Hat Enterprise Linux 7 servers have certain support limits.

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

- ✦ For host systems: <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 <https://access.redhat.com/site/supported-hypervisors>.



Note

To verify whether your processor supports the virtualization extensions and for information on enabling the virtualization extensions if they are disabled, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

Nested Virtualization

As of version 7.2, Red Hat Enterprise Linux offers the nested virtualization feature as a Technology Preview.

Nested virtualization makes it possible to use KVM guests as hosts, and thus enables the user to create one or more KVM guests within each KVM guest. This for example allows debugging hypervisors on a virtual machine or testing larger virtual deployments on a limited amount of physical machines.

3.2. libvirt and libvirt tools

The *libvirt* package is a hypervisor-independent virtualization API that is able to interact with the virtualization capabilities of a range of operating systems.

The *libvirt* package provides:

- ✦ A standard, generic, and stable virtualization layer to securely manage virtual machines on a host.
- ✦ A standard interface for managing local systems and networked hosts.
- ✦ All of the APIs required to provision, create, modify, monitor, control, migrate, and stop virtual machines, but only if the hypervisor supports these operations. Although multiple hosts may be accessed with **libvirt** simultaneously, the APIs are limited to single node operations.

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 can be located on separate physical machines from the host using secure protocols.

Red Hat Enterprise Linux 7 supports **libvirt** and includes **libvirt**-based tools as its default method for virtualization management (as in Red Hat Enterprise 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 Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6, and Red Hat Enterprise Linux 7.

virsh

The **virsh** command-line tool is built on the **libvirt** management API and operates as an alternative to the graphical **virt-manager** application. The **virsh** command can be used in read-only mode by unprivileged users or, with root access, full administration functionality. The **virsh** command is ideal for scripting virtualization administration and provides many functions such as installing, listing, starting, and stopping virtual machines.

virt-manager

virt-manager is a graphical desktop tool for managing virtual machines. It allows access to graphical guest consoles and can be used to perform virtualization administration, virtual machine creation, migration, and configuration tasks. The ability to view virtual machines, host statistics, device information and performance graphs is also provided. The local hypervisor and remote hypervisors can be managed through a single interface.



Note

For more information on **virsh** and **virt-manager**, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

3.3. Virtualized hardware devices

Virtualization on Red Hat Enterprise Linux 7 allows virtual machines to use the host's physical hardware as three distinct types of devices:

- ✦ Virtualized and emulated devices
- ✦ Para-virtualized 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 in software. These emulated hardware devices are crucial for virtualizing operating systems.

Emulated devices are virtual devices which exist entirely in software.

Emulated drivers may use either a physical device or a virtual software device. Emulated drivers are 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, and mouse) and recognized by the Linux kernel, may 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 their use, regardless of the number of host CPUs.

Emulated graphics devices

Two emulated graphics devices are provided. These devices can be connected to with the SPICE (Simple Protocol for Independent Computing Environments) protocol or with VNC:

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

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
- USB 3.0 xHCI host controller (Technology Preview in Red Hat Enterprise Linux 7.2)

Emulated sound devices

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

- Red Hat Enterprise Linux 7, for the x86_64 architecture
- Red Hat Enterprise Linux 6, for i386 and x86_64 architectures
- Red Hat Enterprise Linux 5, for i386 and x86_64 architectures
- Red Hat Enterprise Linux 4, for i386 and x86_64 architectures

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

- **ac97**, an emulated Intel 82801AA AC97 Audio compatible sound card

Emulated watchdog devices

Red Hat Enterprise Linux 6.0 and above provides 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 i386 and x86_64 architectures for guest operating systems Red Hat Enterprise Linux 6.2 and above.

Emulated network devices

There are two emulated network devices available:

- ✦ The **e1000** device emulates an Intel E1000 network adapter (Intel 82540EM, 82573L, 82544GC).
- ✦ The **rtl8139** device emulates a Realtek 8139 network adapter.

Emulated storage drivers

Storage devices and storage pools can use these emulated drivers 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.

The emulated AHCI controller

The emulated Advanced Host Controller Interface (AHCI) bus is an alternative to IDE. It is available as a Technology Preview in Red Hat Enterprise Linux 7.2.



Note

For more information on KVM storage features, refer to [Section 3.4, “Storage”](#).

3.3.2. Para-virtualized devices

Para-virtualization provides a fast and efficient means of communication for guests to use devices on the host machine. KVM provides para-virtualized devices to virtual machines using the Virtio API as a layer between the hypervisor and guest.

Some para-virtualized devices decrease I/O latency and increase I/O throughput to near bare-metal levels, while other para-virtualized devices add functionality to virtual machines that is not otherwise available. It is recommended to use para-virtualized 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. Para-virtualized device drivers allow the guest operating system access to physical devices on the host system.

To use this device, the para-virtualized device drivers must be installed on the guest operating system. By default, the para-virtualized 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.



Note

For more information on using the para-virtualized devices and drivers, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

The para-virtualized network device (virtio-net)

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

The para-virtualized block device (virtio-blk)

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

The para-virtualized controller device (virtio-scsi)

The para-virtualized SCSI controller device 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.

virtio-scsi is fully supported for the following guest operating systems:

- ✦ Red Hat Enterprise Linux 7
- ✦ Red Hat Enterprise Linux 6.4 and above

The para-virtualized 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 para-virtualized clock. Additionally, the para-virtualized clock assists with time adjustments needed after a guest runs S3 or suspend to RAM operations.

The para-virtualized serial device (virtio-serial)

The para-virtualized 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 balloon *inflation*), 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 para-virtualized random number generator (`virtio-rng`)

The para-virtualized random number generator enables virtual machines to collect entropy, or randomness, directly from the host to use for encrypted data and security. Virtual machines can often be starved of entropy because typical inputs (such as hardware usage) are unavailable. Sourcing entropy can be time-consuming; `virtio-rng` makes this process faster by injecting entropy directly into guest virtual machines from the host.

The para-virtualized graphics card (QXL)

The para-virtualized 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*. Device assignment is also known as *passthrough*.

VFIO device assignment

Virtual Function I/O (VFIO) is a new kernel driver in Red Hat Enterprise Linux 7 that provides virtual machines with high performance access to physical hardware.

VFIO attaches PCI devices on the host system directly to virtual machines, allowing guests exclusive access to PCI devices for a range of tasks. This allows PCI devices to appear and behave as if they were physically attached to the guest virtual machine.

VFIO improves on previous PCI device assignment architecture by moving device assignment out of the KVM hypervisor, and enforcing device isolation at the kernel level. VFIO offers better security and is compatible with secure boot. It is the default device assignment mechanism in Red Hat Enterprise Linux 7.

VFIO increases the number of assigned devices to 32 in Red Hat Enterprise Linux 7, from a maximum 8 devices in Red Hat Enterprise Linux 6. VFIO also supports assignment of NVIDIA GPUs.



Note

For more information on VFIO device assignment, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

USB, PCI, and SCSI passthrough

The KVM hypervisor supports attaching USB, PCI, and SCSI devices on the host system to virtual machines. USB, PCI, and SCSI device assignment allows the devices to appear and behave as if they were physically attached to the virtual machine. Thus, it allows guests to have exclusive access to these devices for a variety of tasks.



Note

For more information on USB, PCI, and SCSI passthrough, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide*.

SR-IOV

SR-IOV (Single Root I/O Virtualization) is a PCI Express (PCI-e) 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 via 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 7 Virtualization Deployment and Administration 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.

3.3.4. Guest CPU models

CPU models define which host CPU features are exposed to the guest operating system. **qemu-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 **qemu-kvm** code.

To allow safe migration of virtual machines between hosts with different sets of CPU features, **qemu-kvm** does not expose all features from the host CPU to guest operating systems by default. Instead, CPU features are exposed based on the chosen CPU model. If a virtual machine has a given CPU feature enabled, it is not possible to migrate it 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 7 Virtualization Deployment and Administration Guide*.

3.4. Storage

Storage for virtual machines is abstracted from the physical storage allocated to the virtual machine. It is attached to the virtual machine using the para-virtualized or emulated block device drivers.

3.4.1. 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 7 Virtualization Deployment and Administration Guide* for more information.

Local storage pools

Local storage pools are directly attached to the host server. They include local directories, directly attached disks, physical partitions, and 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 as 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**.

3.4.2. Storage volumes

Storage pools are further 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.

3.4.3. Emulated storage devices

Virtual machines can be presented with a range of storage devices that are emulated by the host. Each type of storage device is appropriate for specific use cases, allowing for maximum flexibility and compatibility with guest operating systems.

virtio-scsi

virtio-scsi is the recommended para-virtualized device for guests using large numbers of disks or advanced storage features such as TRIM. Guest driver installation may be necessary on guests using operating systems other than Red Hat Enterprise Linux 7.

virtio-blk

virtio-blk is a para-virtualized 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 allowing the guests to use ISO files or the host CD-ROM drive. 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

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

AHCI

The emulated Advanced Host Controller Interface (AHCI) bus is an alternative to IDE which provides increased features and performance, including communication with Serial ATA (SATA) devices.

AHCI is available as a Technology Preview in Red Hat Enterprise Linux 7.2.

3.4.4. Host storage

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

Image files

Image files are 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.

Red Hat Enterprise Linux 7.0 and newer supports the qcow2 version 3 image file format.

LVM volumes

Logical volumes 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 and raw disks or LUNs can be presented to the guest. This allows SAN or iSCSI LUNs as well as local CD-ROM media to be used by the guest 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 allows high-performance clustered storage over the network.

Red Hat Enterprise Linux 7 includes native support for disk images on 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 7 Virtualization Deployment and Administration Guide*.

Chapter 4. Virtualization Tools

This chapter provides an introduction to the many tools available to assist with virtualization.

4.1. `virsh`

`virsh` is a command-line interface (CLI) tool for managing the hypervisor and guest virtual machines. The `virsh` command-line tool is built on the `libvirt` management API and operates as an alternative to the `qemu-kvm` command and the graphical `virt-manager` application. The `virsh` command can be used in read-only mode by unprivileged users or, with root access, full administrative functionality. The `virsh` command is ideal for scripting virtualization administration. In addition the `virsh` tool is a 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 7 Virtualization Deployment and Administration Guide* for more information about managing virtual machines with `virsh`.

4.2. `virt-manager`

`virt-manager` is a lightweight graphical tool for managing virtual machines. 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. This tool is provided by its own package called `virt-manager`.



Note

Refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide* for more information about managing virtual machines with `virt-manager`.

4.3. `virt-install`

`virt-install` is a command-line tool to provision new virtual machines. It supports both text-based and graphical installations, using serial console, 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 use the kickstart method to prepare the guest, allowing for easy automation of installation. This tool is included in the `virt-install` package.



Note

Refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide* for more information about using `virt-install`.

4.4. `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 can 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 7 Virtualization Deployment and Administration Guide* for more information about `guestfish`.

4.5. GNOME Boxes

`Boxes` is a lightweight graphical desktop virtualization tool used to view and access virtual machines and remote systems. `Boxes` provides a way to test different operating systems and applications from the desktop with minimal configuration. Virtual systems can be installed manually or with the express installation function, which automatically pre-configures virtual machines with optimal settings. This tool is provided by its own package called `gnome-boxes`.



Note

Refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide* for more information about using **GNOME Boxes**.

4.6. Other useful tools

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

`guestmount`

A command-line tool used to mount virtual machine file systems and disk images on the host machine. This tool is installed as part of the `libguestfs-tools-c` package.



Warning

Using **guestmount** in `--rw` (read/write) mode to access a disk that is currently being used by a guest can cause the disk to become corrupted. Do not use **guestmount** in `--rw` (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-builder

A command-line tool for quickly building and customizing new virtual machines. This tool is installed in Red Hat Enterprise Linux 7.1 and later as part of the *libguestfs-tools* package.

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* package.

virt-customize

A command-line tool for customizing virtual machine disk images. `virt-customize` can be used to install packages, edit configuration files, run scripts, and set passwords. This tool is installed in Red Hat Enterprise Linux 7.1 and later as part of the *libguestfs-tools* 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* package.

virt-diff

A command-line tool for showing differences between the file systems of two virtual machines, for example, to discover which files have changed between snapshots. This tool is installed in Red Hat Enterprise Linux 7.1 and later as part of the *libguestfs-tools* 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* package.



Warning

Using **virt-edit** on live virtual machines can 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* package.

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* package.

virt-log

A command-line tool for listing log files from virtual machines. This tool is installed in Red Hat Enterprise Linux 7.1 and later as part of the *libguestfs-tools* 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* 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-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* package.



Warning

Using **virt-rescue** on running virtual machines can 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* package.



Important

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

virt-sparsify

A command-line tool to make a virtual machine disk (or any disk image) thin-provisioned. The tool can convert free space in the disk image to free space on the host.

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* 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-in

A command-line archive tool for unpacking an uncompressed tarball into a virtual machine disk image or specified libvirt domain. This tool is installed as part of the *libguestfs-tools* package.



Warning

Using **virt-tar-in** command on a live virtual machine can cause disk corruption in the virtual machine. The virtual machine must be shut down before using this command.

virt-tar-out

A command-line archive tool for packing a virtual machine disk image directory into a tarball. This tool is installed as part of the *libguestfs-tools* package.

virt-top

A command-line utility similar to **top**, which shows statistics related to virtualized domains. This tool is contained in its own package: *virt-top*.

virt-v2v

A command-line tool for converting virtual machines from a foreign hypervisor to run on KVM managed by libvirt, on Red Hat Enterprise Linux OpenStack Platform, or on Red Hat Enterprise Virtualization. Currently, *virt-v2v* can convert Red Hat Enterprise Linux and Windows guests running on Xen and VMware ESX. The *virt-v2v* tool is installed in Red Hat Enterprise Linux 7.1 and later as part of the *libguestfs* package.

virt-viewer

A minimal tool for displaying the graphical console of a virtual machine via the VNC and SPICE protocols. This tool is provided by its own package: *virt-viewer*.

virt-what

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

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 its own package: *virt-who*.

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.

Chapter 5. Quick Start Guide to Virtualization in Red Hat Enterprise Linux 7

This chapter provides instructions for installing the basic form of virtualization in Red Hat Enterprise Linux 7.

The chapter begins with an outline of the minimum system specifications and required packages to use virtualization. After these basic packages are installed, follow the procedure in [Section 5.2.2, “Creating a virtual machine with Virtual Machine Manager”](#) to create a basic virtual machine on **Virtual Machine Manager**.



Note

This tutorial uses Virtual Machine Manager to quickly create a virtual machine to try out KVM virtualization. To set up virtual machines with the capabilities necessary for a production environment, refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide* for detailed information about system requirements and recommendations for running virtual machines.

5.1. Basic requirements and setup

To set up a KVM virtual machine on Red Hat Enterprise 7, your system must meet several criteria:

Architecture

On certain architectures, KVM virtualization is not supported with Red Hat Enterprise Linux 7. Notably, these include IBM Power Systems.

Disk space and RAM

For virtualization to work on your system, you require at minimum:

- 6 GB free disk space;
- 2 GB of RAM.

Required packages for virtualization

Before you can use virtualization, a basic set of virtualization packages must be installed on your computer. The virtualization packages can be installed either during the host installation sequence, or after host installation using the **yum** command and the Red Hat Customer Portal.

To install the packages, your host machine must be registered and subscribed to the Red Hat Customer Portal. To register via Red Hat Subscription Manager, 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>.

Procedure 5.1. Installing the virtualization packages with yum

To use virtualization on Red Hat Enterprise Linux, you require at minimum the **qemu-kvm**, **qemu-img**, and **libvirt** packages. These provide the user-level KVM emulator, disk image manager, and virtualization management tools on the host Red Hat Enterprise Linux system

1. Install the *qemu-kvm* and *qemu-img* packages with the following command:

```
# yum install qemu-kvm qemu-img
```

2. Install the *virt-manager* package with the following command:

```
# yum install virt-manager
```

3. Additionally, for this tutorial, download a Red Hat Enterprise Linux 7 ISO image to your system to create the guest virtual machine's operating system.

Red Hat Enterprise Linux ISO files can be found on the Red Hat Customer Portal at <https://access.redhat.com/downloads/>.

5.2. Deploying a virtual machine with Virtual Machine Manager

The **Virtual Machine Manager**, also known as **virt-manager**, is a graphical tool for quickly deploying 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.

5.2.1. Introduction to Virtual Machine Manager

Open the **Virtual Machine Manager** application from the **Applications** menu and **System Tools** submenu.

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




Figure 5.1. The Virtual Machine Manager interface

Commonly-used functions include:

- ✦ **1 Create new virtual machine:** Start here to create a new virtual machine.
- ✦ **2 Virtual machines:** A list of all virtual machines, or guests. 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 under **CPU usage**.

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 dropdown 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.

5.2.2. 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 5.2. Creating a guest virtual machine with Virtual Machine Manager

1. Open Virtual Machine Manager

Launch the **Virtual Machine Manager** application from the **Applications** menu and **System Tools** submenu.

2. Create a new virtual machine

Click the **Create a new virtual machine** button ([Figure 5.2, “The Create a new virtual machine button”](#)) to open the **New VM** wizard.




Figure 5.2. The Create a new virtual machine button

3. Specify the installation method

Start the creation process by choosing the method of installing the new virtual machine.

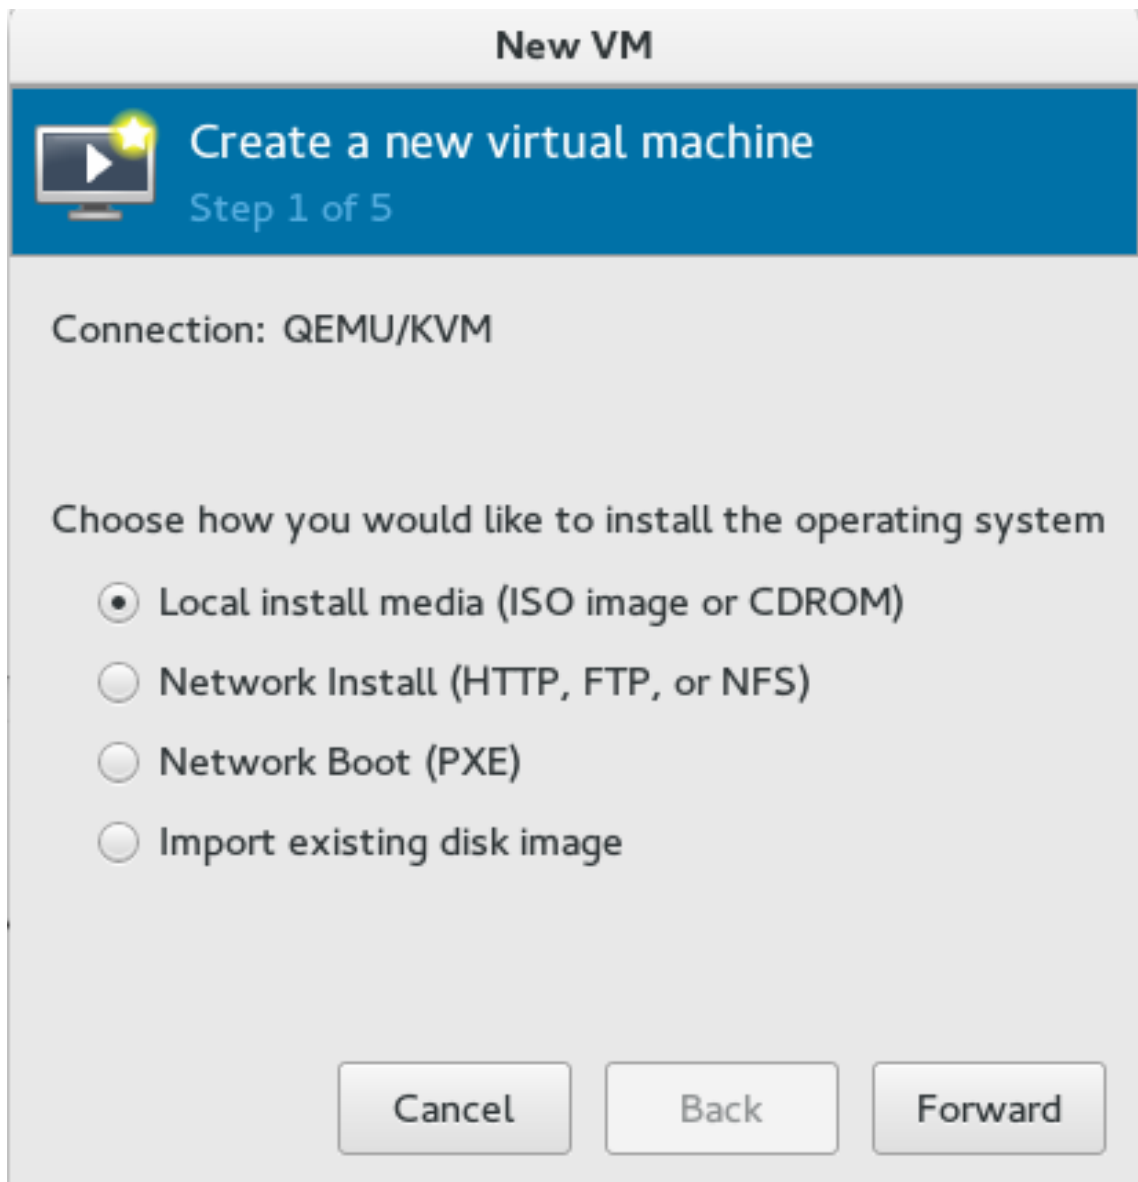


Figure 5.3. Select the installation method

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

4. Locate installation media

Provide the location of the ISO that the guest virtual machine will install the operating system from.

For this example, locate the ISO downloaded in [Procedure 5.1, “Installing the virtualization packages with yum”](#) on your machine by using the **Browse** to provide the existing storage path.

Select the **OS type** and **Version** of the installation, ensuring that you select the appropriate OS type for your virtual machine. For this example, select **Linux** from the dropdown in **OS type** and choose **Red Hat Enterprise Linux 7** from the dropdown in **Version**, and click the **Forward** button. Note that you can also let the Virtual Machine Manager automatically detect the operating system by clicking the respective checkbox.

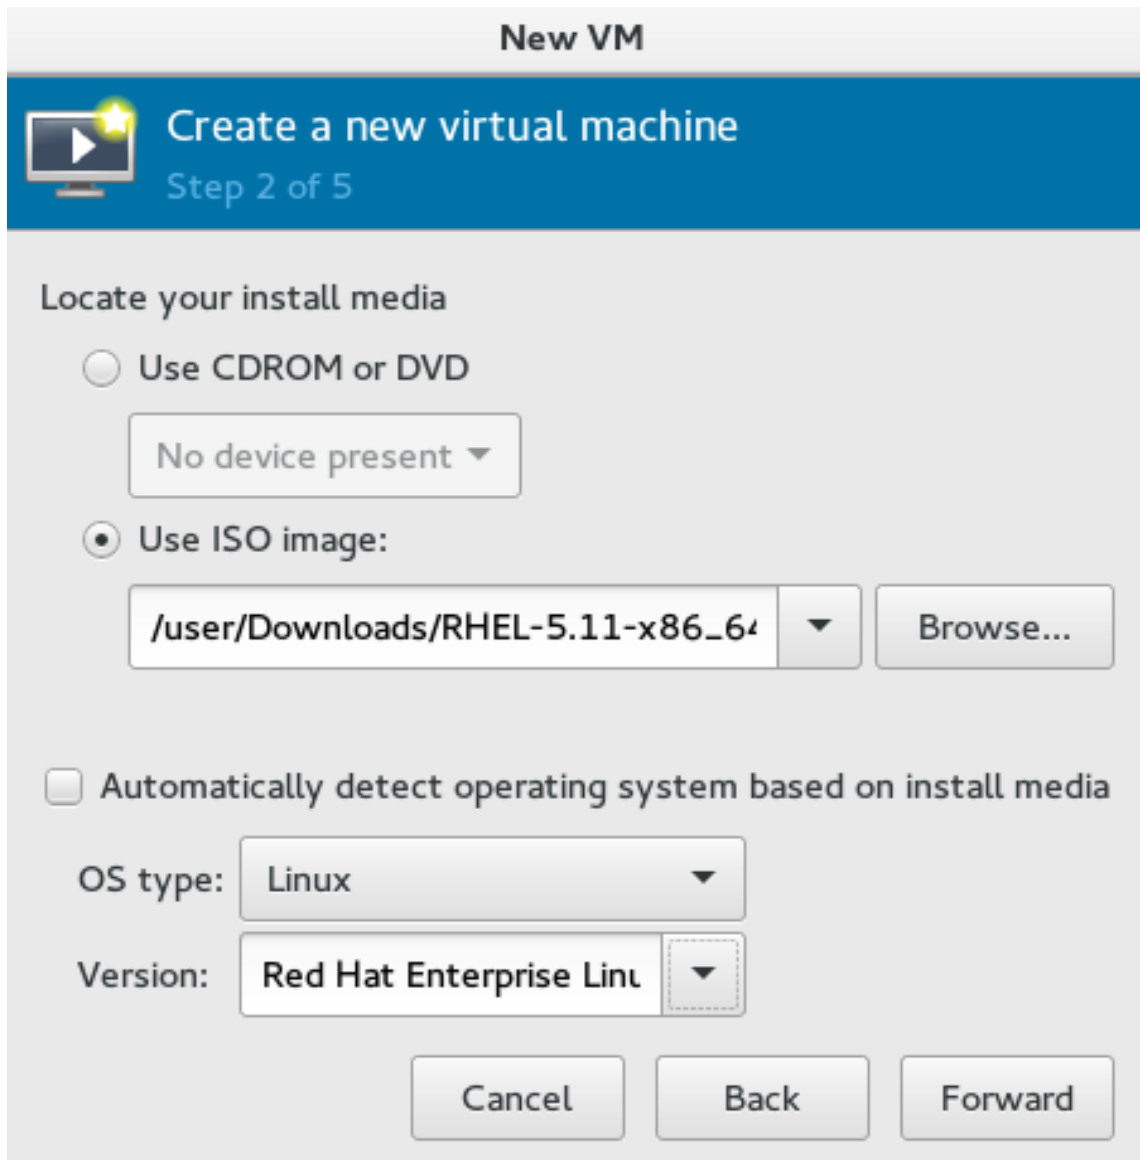


Figure 5.4. Local ISO image installation

5. Configure memory and CPU

Select 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, select the default settings and click the **Forward** button.




Figure 5.5. Configuring CPU and Memory

6. Configure storage

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, select the default settings and click the **Forward** button.




Figure 5.6. Configuring storage

7. Naming and final overview

Finally, select a name for the virtual machine, verify the settings, and click the **Finish** button. **Virtual Machine Manager** will then create a virtual machine with your specified hardware settings.

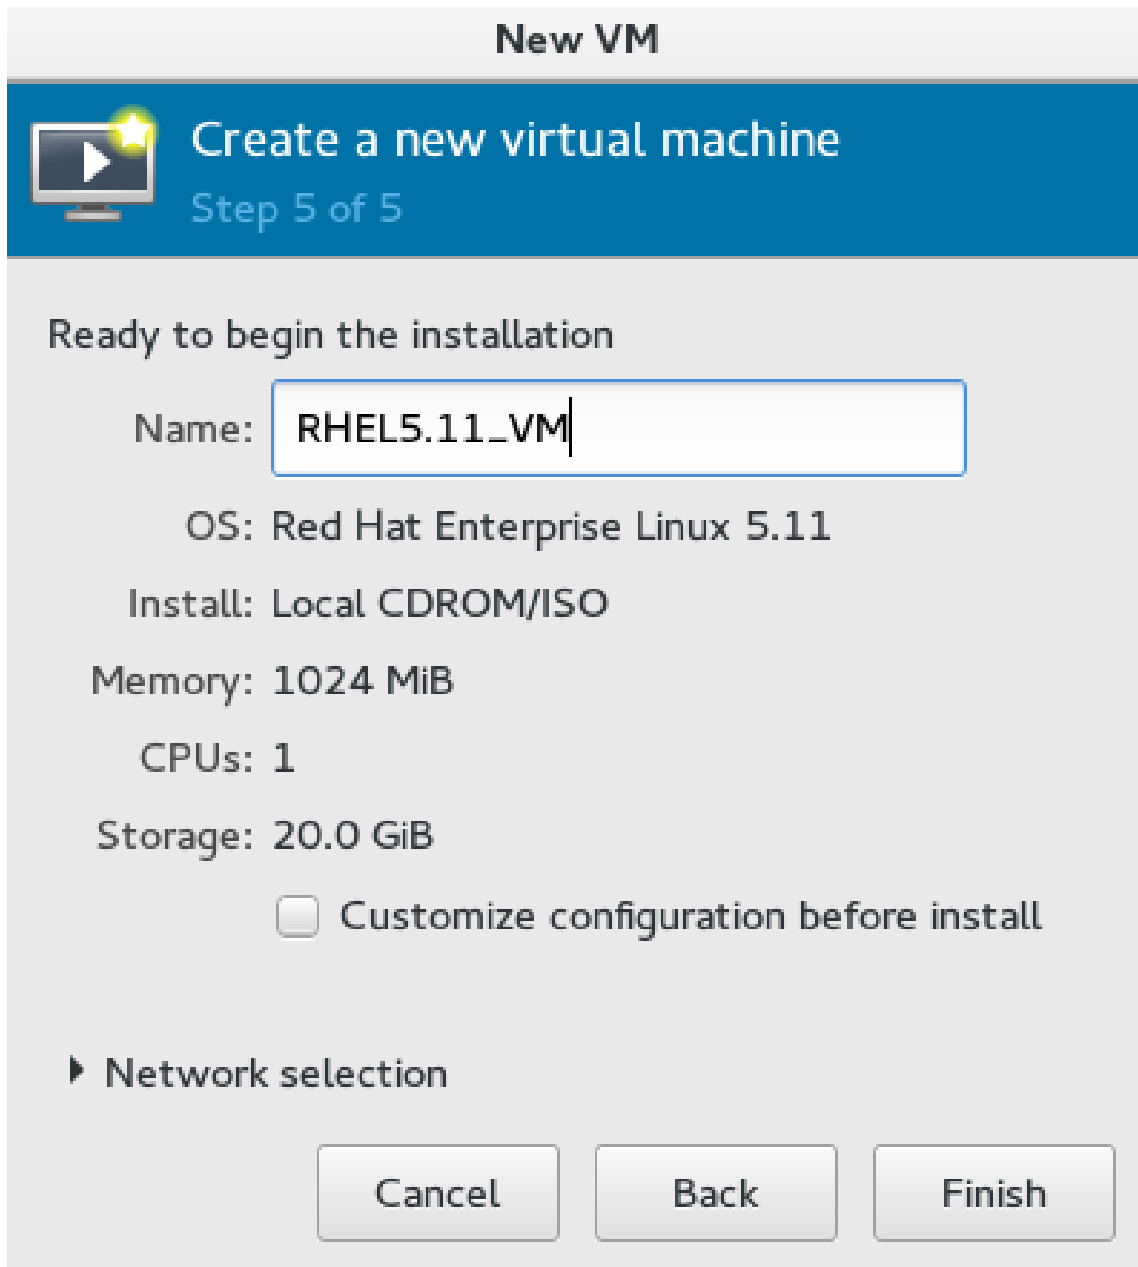


Figure 5.7. Naming and verification

After **Virtual Machine Manager** has created your Red Hat Enterprise Linux 7 virtual machine, 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*.

5.2.3. 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 5.8. The guest virtual machine console

- 1
 ✦ **Show the graphical console:** This shows what is on the virtual machine's screen. The virtual machine can be operated the same as a physical machine from the console.
- 2
 ✦ **Show virtual hardware details:** This window shows details of the virtual hardware that the guest is using. It provides an overview of basic system details, performance, processor, memory, and boot settings, plus 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 dropdown 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

Refer to the *Red Hat Enterprise Linux 7 Virtualization Deployment and Administration Guide* for more information about using the **Virtual Machine Manager** to create and run virtual machines.

Appendix A. Revision History

Revision 1.0-45	Thu Nov 12 2015	Laura Novich
Preparing document for 7.2 GA publication.		
Revision 1.0-44	Thu Nov 12 2015	Laura Novich
Preparing document for 7.2 GA publication.		
Revision 1.0-43	Thu Oct 08 2015	Jiri Herrmann
Cleaned up the Revision History		
Revision 1.0-42	Sun Jun 28 2015	Jiri Herrmann
Updated for the 7.2 beta release		