



Red Hat Virtualization 4.4

Virtual Machine Management Guide

Managing virtual machines in Red Hat Virtualization

Red Hat Virtualization 4.4 Virtual Machine Management Guide

Managing virtual machines in Red Hat Virtualization

Red Hat Virtualization Documentation Team

Red Hat Customer Content Services

rhev-docs@redhat.com

Legal Notice

Copyright © 2024 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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, the Red Hat 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 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

This document describes the installation, configuration, and administration of virtual machines in Red Hat Virtualization.

Table of Contents

CHAPTER 1. INTRODUCTION	4
1.1. AUDIENCE	4
1.2. SUPPORTED VIRTUAL MACHINE OPERATING SYSTEMS	4
1.3. VIRTUAL MACHINE PERFORMANCE PARAMETERS	4
1.4. INSTALLING SUPPORTING COMPONENTS ON CLIENT MACHINES	4
CHAPTER 2. INSTALLING RED HAT ENTERPRISE LINUX VIRTUAL MACHINES	7
2.1. CREATING A VIRTUAL MACHINE	7
2.2. STARTING THE VIRTUAL MACHINE	9
2.3. ENABLING THE REQUIRED REPOSITORIES	13
2.4. INSTALLING GUEST AGENTS AND DRIVERS	14
CHAPTER 3. INSTALLING WINDOWS VIRTUAL MACHINES	18
3.1. CREATING A VIRTUAL MACHINE	18
3.2. STARTING THE VIRTUAL MACHINE USING RUN ONCE	19
3.3. INSTALLING GUEST AGENTS AND DRIVERS	21
CHAPTER 4. ADDITIONAL CONFIGURATION	27
4.1. CONFIGURING OPERATING SYSTEMS WITH OSINFO	27
4.2. CONFIGURING SINGLE SIGN-ON FOR VIRTUAL MACHINES	27
4.3. CONFIGURING USB DEVICES	30
4.4. CONFIGURING MULTIPLE MONITORS	31
4.5. CONFIGURING CONSOLE OPTIONS	32
4.6. CONFIGURING A WATCHDOG	37
4.7. CONFIGURING VIRTUAL NUMA	42
4.8. CONFIGURING SATELLITE ERRATA VIEWING FOR A VIRTUAL MACHINE	44
4.9. CONFIGURING HEADLESS VIRTUAL MACHINES	45
4.10. CONFIGURING HIGH PERFORMANCE VIRTUAL MACHINES, TEMPLATES, AND POOLS	46
4.11. CONFIGURING THE TIME ZONE	53
CHAPTER 5. EDITING VIRTUAL MACHINES	55
5.1. EDITING VIRTUAL MACHINE PROPERTIES	55
5.2. NETWORK INTERFACES	56
5.3. VIRTUAL DISKS	58
5.4. VIRTUAL MEMORY	62
5.5. HOT PLUGGING VCPUS	63
5.6. PINNING A VIRTUAL MACHINE TO MULTIPLE HOSTS	65
5.7. VIEWING VIRTUAL MACHINES PINNED TO A HOST	65
5.8. CHANGING THE CD FOR A VIRTUAL MACHINE	66
5.9. SMART CARD AUTHENTICATION	66
CHAPTER 6. ADMINISTRATIVE TASKS	68
6.1. SHUTTING DOWN A VIRTUAL MACHINE	68
6.2. SUSPENDING A VIRTUAL MACHINE	68
6.3. REBOOTING OR RESETTING A VIRTUAL MACHINE	68
6.4. REMOVING A VIRTUAL MACHINE	69
6.5. CLONING A VIRTUAL MACHINE	70
6.6. UPDATING VIRTUAL MACHINE GUEST AGENTS AND DRIVERS	70
6.7. VIEWING RED HAT SATELLITE ERRATA FOR A VIRTUAL MACHINE	72
6.8. VIRTUAL MACHINES AND PERMISSIONS	72
6.9. SNAPSHOTS	76
6.10. HOST DEVICES	78

6.11. AFFINITY GROUPS	82
6.12. AFFINITY LABELS	87
6.13. EXPORTING AND IMPORTING VIRTUAL MACHINES AND TEMPLATES	89
6.14. MIGRATING VIRTUAL MACHINES BETWEEN HOSTS	104
6.15. IMPROVING UPTIME WITH VIRTUAL MACHINE HIGH AVAILABILITY	111
6.16. OTHER VIRTUAL MACHINE TASKS	113
CHAPTER 7. TEMPLATES	120
7.1. ABOUT TEMPLATES	120
7.2. SEALING VIRTUAL MACHINES IN PREPARATION FOR DEPLOYMENT AS TEMPLATES	120
7.3. CREATING A TEMPLATE	122
7.4. EDITING A TEMPLATE	123
7.5. DELETING A TEMPLATE	124
7.6. EXPORTING TEMPLATES	124
7.7. IMPORTING TEMPLATES	125
7.8. TEMPLATES AND PERMISSIONS	126
7.9. USING CLOUD-INIT TO AUTOMATE THE CONFIGURATION OF VIRTUAL MACHINES	128
7.10. USING SYSPREP TO AUTOMATE THE CONFIGURATION OF VIRTUAL MACHINES	132
7.11. CREATING A VIRTUAL MACHINE BASED ON A TEMPLATE	134
7.12. CREATING A CLONED VIRTUAL MACHINE BASED ON A TEMPLATE	135
APPENDIX A. REFERENCE: SETTINGS IN ADMINISTRATION PORTAL AND VM PORTAL WINDOWS	137
A.1. EXPLANATION OF SETTINGS IN THE NEW VIRTUAL MACHINE AND EDIT VIRTUAL MACHINE WINDOWS	137
A.2. EXPLANATION OF SETTINGS IN THE RUN ONCE WINDOW	167
A.3. EXPLANATION OF SETTINGS IN THE NEW NETWORK INTERFACE AND EDIT NETWORK INTERFACE WINDOWS	173
A.4. EXPLANATION OF SETTINGS IN THE NEW VIRTUAL DISK AND EDIT VIRTUAL DISK WINDOWS	174
A.5. EXPLANATION OF SETTINGS IN THE NEW TEMPLATE WINDOW	178
APPENDIX B. VIRT-SYSPREP OPERATIONS	181
APPENDIX C. LEGAL NOTICE	182

CHAPTER 1. INTRODUCTION

A virtual machine is a software implementation of a computer. The Red Hat Virtualization environment enables you to create virtual desktops and virtual servers.

Virtual machines consolidate computing tasks and workloads. In traditional computing environments, workloads usually run on individually administered and upgraded servers. Virtual machines reduce the amount of hardware and administration required to run the same computing tasks and workloads.

1.1. AUDIENCE

Most virtual machine tasks in Red Hat Virtualization can be performed in both the VM Portal and Administration Portal. However, the user interface differs between each portal, and some administrative tasks require access to the Administration Portal. Tasks that can only be performed in the Administration Portal will be described as such in this book. Which portal you use, and which tasks you can perform in each portal, is determined by your level of permissions. Virtual machine permissions are explained in [Virtual Machines and Permissions](#).

The VM Portal's user interface is described in the [Introduction to the VM Portal](#).

The Administration Portal's user interface is described in the [Administration Guide](#).

The creation and management of virtual machines through the Red Hat Virtualization REST API is documented in the [REST API Guide](#).

1.2. SUPPORTED VIRTUAL MACHINE OPERATING SYSTEMS

See [Certified Guest Operating Systems in Red Hat OpenStack Platform, Red Hat Virtualization and OpenShift Virtualization](#) for a current list of supported operating systems.

For information on customizing the operating systems, see [Configuring operating systems with osinfo](#).

1.3. VIRTUAL MACHINE PERFORMANCE PARAMETERS

For information on the parameters that Red Hat Virtualization virtual machines can support, see [Red Hat Enterprise Linux technology capabilities and limits](#) and [Virtualization limits for Red Hat Virtualization](#).

1.4. INSTALLING SUPPORTING COMPONENTS ON CLIENT MACHINES

1.4.1. Installing Console Components

A console is a graphical window that allows you to view the start up screen, shut down screen, and desktop of a virtual machine, and to interact with that virtual machine in a similar way to a physical machine. In Red Hat Virtualization, the default application for opening a console to a virtual machine is Remote Viewer, which must be installed on the client machine prior to use.

1.4.1.1. Installing Remote Viewer on Red Hat Enterprise Linux

The Remote Viewer application provides users with a graphical console for connecting to virtual machines. Once installed, it is called automatically when attempting to open a SPICE session with a virtual machine. Alternatively, it can also be used as a standalone application. Remote Viewer is included

in the **virt-viewer** package provided by the base Red Hat Enterprise Linux Workstation and Red Hat Enterprise Linux Server repositories.

Procedure

1. Install the **virt-viewer** package:

```
# dnf install virt-viewer
```

2. Restart your browser for the changes to take effect.

You can now connect to your virtual machines using either the SPICE protocol or the VNC protocol.

1.4.1.2. Installing Remote Viewer on Windows

The Remote Viewer application provides users with a graphical console for connecting to virtual machines. Once installed, it is called automatically when attempting to open a SPICE session with a virtual machine. Alternatively, it can also be used as a standalone application.

Installing Remote Viewer on Windows

1. Open a web browser and download one of the following installers according to the architecture of your system.

- Virt Viewer for 32-bit Windows:

```
https://your-manager-fqdn/ovirt-engine/services/files/spice/virt-viewer-x86.msi
```

- Virt Viewer for 64-bit Windows:

```
https://your-manager-fqdn/ovirt-engine/services/files/spice/virt-viewer-x64.msi
```

2. Open the folder where the file was saved.
3. Double-click the file.
4. Click **Run** if prompted by a security warning.
5. Click **Yes** if prompted by User Account Control.

Remote Viewer is installed and can be accessed via **Remote Viewer** in the **VirtViewer** folder of **All Programs** in the start menu.

1.4.1.3. Installing usbdk on Windows

usbdk is a driver that enables **remote-viewer** exclusive access to USB devices on Windows operating systems. Installing **usbdk** requires Administrator privileges. Note that the previously supported **USB Clerk** option has been deprecated and is no longer supported.

Installing usbdk on Windows

1. Open a web browser and download one of the following installers according to the architecture of your system.

- **usbdk** for 32-bit Windows:

https://[*your manager's address*]/ovirt-engine/services/files/spice/usbdk-x86.msi

- **usbdk** for 64-bit Windows:

https://[*your manager's address*]/ovirt-engine/services/files/spice/usbdk-x64.msi

2. Open the folder where the file was saved.
3. Double-click the file.
4. Click **Run** if prompted by a security warning.
5. Click **Yes** if prompted by User Account Control.

CHAPTER 2. INSTALLING RED HAT ENTERPRISE LINUX VIRTUAL MACHINES

Installing a Red Hat Enterprise Linux virtual machine involves the following key steps:

1. [Create a virtual machine](#). You must add a virtual disk for storage, and a network interface to connect the virtual machine to the network.
2. [Start the virtual machine](#) and install an operating system. See your operating system's documentation for instructions.
 - Red Hat Enterprise Linux 6: [Installing Red Hat Enterprise Linux 6.9 for all architectures](#)
 - Red Hat Enterprise Linux 7: [Installing Red Hat Enterprise Linux 7 on all architectures](#)
 - Red Hat Enterprise Linux Atomic Host 7: [Red Hat Enterprise Linux Atomic Host 7 Installation and Configuration Guide](#)
 - Red Hat Enterprise Linux 8: [Installing Red Hat Enterprise Linux 8 using the graphical user interface](#)
3. [Enable the required repositories](#) for your operating system.
4. [Install guest agents and drivers](#) for additional virtual machine functionality.

2.1. CREATING A VIRTUAL MACHINE

When creating a new virtual machine, you specify its settings. You can edit some of these settings later, including the chipset and BIOS type. For more information, see [UEFI and the Q35 chipset](#) in the *Administration Guide*. .Prerequisites




NOTE

Before you can use this virtual machine, you must:

- Install an operating system
 - Use a pre-installed image by [Creating a Cloned Virtual Machine Based on a Template](#)
 - Use a pre-installed image from an attached pre-installed Disk
 - Install an operating system through the PXE boot menu or from an ISO file
- Register with the Content Delivery Network

Procedure

1. Click **Compute** → **Virtual Machines**.
2. Click **New**. This opens the **New Virtual Machine** window.
3. Select an **Operating System** from the drop-down list.

**NOTE**


If you selected **Red Hat Enterprise Linux CoreOS** as the operating system, you may need to set the initialization method by configuring **Ignition** settings in the **Advanced Options Initial Run** tab. See [Configuring Ignition](#).

4. Enter a **Name** for the virtual machine.
5. Add storage to the virtual machine: under **Instance Images**, click **Attach** or **Create** to select or create a virtual disk .
 - Click **Attach** and select an existing virtual disk.
or
 - Click **Create** and enter a **Size(GB)** and **Alias** for a new virtual disk. You can accept the default settings for all other fields, or change them if required. See [Explanation of settings in the New Virtual Disk and Edit Virtual Disk windows](#) for more details on the fields for all disk types.
6. Connect the virtual machine to the network. Add a network interface by selecting a vNIC profile from the **nic1** drop-down list at the bottom of the **General** tab.
7. Specify the virtual machine's **Memory Size** on the **System** tab.
8. In the **Boot Options** tab, choose the **First Device** that the virtual machine will use to boot.
9. You can accept the default settings for all other fields, or change them if required. For more details on all fields in the **New Virtual Machine** window, see [Explanation of settings in the New Virtual Machine and Edit Virtual Machine Windows](#).
10. Click **OK**.

The new virtual machine is created and displays in the list of virtual machines with a status of **Down**.

Configuring Ignition


Ignition is the utility that is used by Red Hat Enterprise Linux CoreOS to manipulate disks during initial configuration. It completes common disk tasks, including partitioning disks, formatting partitions, writing files, and configuring users. On first boot, Ignition reads its configuration from the installation media or the location that you specify and applies the configuration to the machines.

**IMPORTANT**

Once Ignition has been configured as the initialization method, it cannot be reversed or re-configured.

1. In the **Add Virtual Machine** or **Edit Virtual Machine** screen, click **Show Advanced Options**.
2. In the **Initial Run** tab, select the **Ignition 2.3.0** option and enter the **VM Hostname**.
3. Expand the **Authorization** option, enter a hashed (SHA-512) password, and enter the password again to verify.
4. If you are using SSH keys for authorization, enter them in the space provided.
5. You can also enter a custom Ignition script in JSON format in the **Ignition Script** field. This

script will run on the virtual machine when it starts. The scripts you enter in this field are custom JSON sections that are added to those produced by the Manager, and allow you to use custom Ignition instructions.



NOTE

If the Red Hat Enterprise Linux CoreOS image you are using contains an Ignition version different than 2.3.0, you need to use a script in the **Ignition Script** field to enforce the Ignition version included in your Red Hat Enterprise Linux CoreOS image.

When you use an Ignition script, the script instructions take precedence over and override any conflicting Ignition settings you configured in the UI.

2.2. STARTING THE VIRTUAL MACHINE

2.2.1. Starting a Virtual Machine

Procedure

1. Click **Compute** → **Virtual Machines** and select a virtual machine with a status of **Down**.
2. Click **Run**.

The **Status** of the virtual machine changes to **Up**, and the operating system installation begins. Open a console to the virtual machine if one does not open automatically.



NOTE

A virtual machine will not start on a host with an overloaded CPU. By default, a host's CPU is considered overloaded if it has a load of more than 80% for 5 minutes, but these values can be changed using scheduling policies. See [Scheduling Policies](#) in the *Administration Guide* for more information.

Troubleshooting

Scenario - the virtual machine fails to boot with the following error message:

Boot failed: not a bootable disk - No Bootable device

Possible solutions to this problem:

- Make sure that **hard disk** is selected in the boot sequence, and the disk that the virtual machine is booting from must be set as **Bootable**.
- Create a [Cloned Virtual Machine Based on a Template](#) .
- Create a new virtual machine with a local boot disk managed by RHV that contains the OS and application binaries.
- Install the OS by booting from the **Network (PXE)** boot option.

Scenario - the virtual machine on IBM POWER9 fails to boot with the following error message:


qemu-kvm: Requested count cache flush assist capability level not supported by kvm, try appending -machine cap-ccf-assist=off

Default risk level protections can prevent VMs from starting on IBM POWER9. To resolve this issue:

1. Create or edit the `/var/lib/obmc/cfam_overrides` on the BMC.
2. Set the firmware risk level to **0**:

```
# Control speculative execution mode
0 0x283a 0x00000000 # bits 28:31 are used for init level -- in this case 0 Kernel and User
protection (safest, default)
0 0x283F 0x20000000 # Indicate override register is valid
```

3. Reboot the host system for the changes to take affect.



NOTE

Overriding the risk level can cause unexpected behavior when running virtual machines.

2.2.2. Opening a console to a virtual machine

Use Remote Viewer to connect to a virtual machine.



NOTE

To allow other users to connect to the VM, make sure you shutdown and restart the virtual machine when you are finished using the console. Alternatively, the administrator can **Disable strict user checking** to eliminate the need for reboot between users. See [Virtual Machine Console Settings Explained](#) for more information.

Procedure

1. Install Remote Viewer if it is not already installed. See [Installing Console Components](#).
2. Click **Compute** → **Virtual Machines** and select a virtual machine.
3. Click **Console**. By default, the browser prompts you to download a file named **console.vv**. When you click to open the file, a console window opens for the virtual machine. You can configure your browser to automatically open these files, such that clicking **Console** simply opens the console.



NOTE

console.vv expires after 120 seconds. If more than 120 seconds elapse between the time the file is downloaded and the time that you open the file, click **Console** again.

Additional resources

- [Automatically connecting to a Virtual Machine](#)
- [Configuring Console Options](#)

2.2.3. Opening a Serial Console to a Virtual Machine

You can access a virtual machine's serial console from the command line instead of opening a console from the Administration Portal or the VM Portal. The serial console is emulated through VirtIO channels, using SSH and key pairs. The Manager acts as a proxy for the connection, provides information about virtual machine placement, and stores the authentication keys. You can add public keys for each user from either the Administration Portal or the VM Portal. You can access serial consoles for only those virtual machines for which you have appropriate permissions.



IMPORTANT

To access the serial console of a virtual machine, the user must have **UserVmManager**, **SuperUser**, or **UserInstanceManager** permission on that virtual machine. These permissions must be explicitly defined for each user. It is not enough to assign these permissions to **Everyone**.

The serial console is accessed through TCP port 2222 on the Manager. This port is opened during **engine-setup** on new installations. To change the port, see [ovirt-vmconsole/README.md](#).

You must configure the following firewall rules to allow a serial console:

- [Rule "M3"](#) for the Manager firewall
- [Rule "H2"](#) for the host firewall


The serial console relies on the **ovirt-vmconsole** package and the **ovirt-vmconsole-proxy** on the Manager and the **ovirt-vmconsole** package and the **ovirt-vmconsole-host** package on the hosts.

These packages are installed by default on new installations. To install the packages on existing installations, [reinstall the hosts](#).

Enabling a Virtual Machine's Serial Console

1. On the virtual machine whose serial console you are accessing, add the following lines to **/etc/default/grub**:

```
GRUB_CMDLINE_LINUX_DEFAULT="console=tty0 console=ttyS0,115200n8"
GRUB_TERMINAL="console serial"
GRUB_SERIAL_COMMAND="serial --speed=115200 --unit=0 --word=8 --parity=no --
stop=1"
```



NOTE

GRUB_CMDLINE_LINUX_DEFAULT applies this configuration only to the default menu entry. Use **GRUB_CMDLINE_LINUX** to apply the configuration to all the menu entries.

If these lines already exist in **/etc/default/grub**, update them. Do not duplicate them.

2. Rebuild **/boot/grub2/grub.cfg**:

- BIOS-based machines:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

- UEFI-based machines:

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

See [GRUB 2 over a Serial Console](#) in the *Red Hat Enterprise Linux 7 System Administrator's Guide* for details.

3. On the client machine from which you are accessing the virtual machine serial console, generate an SSH key pair. The Manager supports standard SSH key types, for example, an RSA key:

```
# ssh-keygen -t rsa -b 2048 -f .ssh/serialconsolekey
```

This command generates a public key and a private key.

4. In the Administration Portal, click **Administration** → **Account Settings** or click the user icon on the header bar and click **Account Settings** to open the **Account Settings** screen.
OR

In the VM Portal, click the Settings icon on the header bar to open the **Account Settings** screen.

5. In the **User's Public Key** text field (Administration Portal) or **SSH Key** field (VM Portal), paste the public key of the client machine that will be used to access the serial console.
6. Click **Compute** → **Virtual Machines** and select a virtual machine.
7. Click **Edit**.
8. In the **Console** tab of the **Edit Virtual Machine** window, select the **Enable VirtIO serial console** check box.

Connecting to a Virtual Machine's Serial Console

On the client machine, connect to the virtual machine's serial console:

- If a single virtual machine is available, this command connects the user to that virtual machine:

```
# ssh -t -p 2222 ovirt-vmconsole@Manager_FQDN -i .ssh/serialconsolekey
Red Hat Enterprise Linux Server release 6.7 (Santiago)
Kernel 2.6.32-573.3.1.el6.x86_64 on an x86_64
USER login:
```

- If more than one virtual machine is available, this command lists the available virtual machines and their IDs:

```
# ssh -t -p 2222 ovirt-vmconsole@Manager_FQDN -i .ssh/serialconsolekey list
1. vm1 [vmid1]
2. vm2 [vmid2]
3. vm3 [vmid3]
> 2
Red Hat Enterprise Linux Server release 6.7 (Santiago)
Kernel 2.6.32-573.3.1.el6.x86_64 on an x86_64
USER login:
```

Enter the number of the machine to which you want to connect, and press **Enter**.

- Alternatively, connect directly to a virtual machine using its unique identifier or its name:

```
# ssh -t -p 2222 ovirt-vmconsole@Manager_FQDN connect --vm-id vmid1
```

```
# ssh -t -p 2222 ovirt-vmconsole@Manager_FQDN connect --vm-name vm1
```

Disconnecting from a Virtual Machine's Serial Console

Press any key followed by ~ . to close a serial console session.

If the serial console session is disconnected abnormally, a TCP timeout occurs. You will be unable to reconnect to the virtual machine's serial console until the timeout period expires.

2.2.4. Automatically Connecting to a Virtual Machine

Once you have logged in, you can automatically connect to a single running virtual machine. This can be configured in the VM Portal.

Procedure

1. In the Virtual Machines page, click the name of the virtual machine to go to the details view.
2. Click the pencil icon beside **Console** and set **Connect automatically** to **ON**.

The next time you log into the VM Portal, if you have only one running virtual machine, you will automatically connect to that machine.

2.3. ENABLING THE REQUIRED REPOSITORIES

To install packages signed by Red Hat you must register the target system to the Content Delivery Network. Then, use an entitlement from your subscription pool and enable the required repositories.

Enabling the Required Repositories Using Subscription Manager

1. Register your system with the Content Delivery Network, entering your Customer Portal user name and password when prompted:

```
# subscription-manager register
```

2. Locate the relevant subscription pools and note down the pool identifiers:

```
# subscription-manager list --available
```

3. Use the pool identifiers to attach the required subscriptions:


```
# subscription-manager attach --pool=pool_id
```

4. When a system is attached to a subscription pool with multiple repositories, only the main repository is enabled by default. Others are available, but disabled. Enable any additional repositories:

```
# subscription-manager repos --enable=repository
```

5. Ensure that all packages currently installed are up to date:

```
# dnf upgrade --nobest
```



NOTE

For versions of Red Hat Enterprise Linux earlier than 8, use the command **yum update** instead of **dnf upgrade**:

```
# yum update
```

See also [Cannot perform yum update on my RHV manager \(ansible conflict\)](#)

2.4. INSTALLING GUEST AGENTS AND DRIVERS

2.4.1. Red Hat Virtualization Guest agents, tools, and drivers

The Red Hat Virtualization guest agents, tools, and drivers provide additional functionality for virtual machines, such as gracefully shutting down or rebooting virtual machines from the VM Portal and Administration Portal. The tools and agents also provide information for virtual machines, including:

- Resource usage
- IP addresses

The guest agents, tools and drivers are distributed as an ISO file that you can attach to virtual machines. This ISO file is packaged as an RPM file that you can install and upgrade from the Manager machine.

You need to install the guest agents and drivers on a virtual machine to enable this functionality for that machine.

Table 2.1. Red Hat Virtualization Guest drivers

Driver	Description	Works on
virtio-net	Paravirtualized network driver provides enhanced performance over emulated devices like rtl.	Server and Desktop.
virtio-block	Paravirtualized HDD driver offers increased I/O performance over emulated devices like IDE by optimizing the coordination and communication between the virtual machine and the hypervisor. The driver complements the software implementation of the virtio-device used by the host to play the role of a hardware device.	Server and Desktop.

Driver	Description	Works on
virtio-scsi	Paravirtualized iSCSI HDD driver offers similar functionality to the virtio-block device, with some additional enhancements. In particular, this driver supports adding hundreds of devices, and names devices using the standard SCSI device naming scheme.	Server and Desktop.
virtio-serial	Virtio-serial provides support for multiple serial ports. The improved performance is used for fast communication between the virtual machine and the host that avoids network complications. This fast communication is required for the guest agents and for other features such as clipboard copy-paste between the virtual machine and the host and logging.	Server and Desktop.
virtio-balloon	Virtio-balloon is used to control the amount of memory a virtual machine actually accesses. It offers improved memory overcommitment.	Server and Desktop.
qxl	A paravirtualized display driver reduces CPU usage on the host and provides better performance through reduced network bandwidth on most workloads.	Server and Desktop.

Table 2.2. Red Hat Virtualization Guest agents and tools

Guest agent/tool	Description	Works on
qemu-guest-agent	Used instead of ovirt-guest-agent-common on Red Hat Enterprise Linux 8 virtual machines. It is installed and enabled by default.	Server and Desktop.

Guest agent/tool	Description	Works on
spice-agent	The SPICE agent supports multiple monitors and is responsible for client-mouse-mode support to provide a better user experience and improved responsiveness than the QEMU emulation. Cursor capture is not needed in client-mouse-mode. The SPICE agent reduces bandwidth usage when used over a wide area network by reducing the display level, including color depth, disabling wallpaper, font smoothing, and animation. The SPICE agent enables clipboard support allowing cut and paste operations for both text and images between client and virtual machine, and automatic guest display setting according to client-side settings. On Windows-based virtual machines, the SPICE agent consists of vdservice and vdagent.	Server and Desktop.

2.4.2. Installing the Guest Agents and Drivers on Red Hat Enterprise Linux

The Red Hat Virtualization guest agents and drivers are provided by the Red Hat Virtualization Agent repository.



NOTE

Red Hat Enterprise Linux 8 virtual machines use the **qemu-guest-agent** service, which is installed and enabled by default, instead of the **ovirt-guest-agent** service. If you need to manually install the guest agent on RHEL 8, follow the procedure below.

Procedure

1. Log in to the Red Hat Enterprise Linux virtual machine.
2. Enable the Red Hat Virtualization Agent repository:

- For Red Hat Enterprise Linux 6

```
# subscription-manager repos --enable=rhel-6-server-rhv-4-agent-rpms
```

- For Red Hat Enterprise Linux 7

```
# subscription-manager repos --enable=rhel-7-server-rh-common-rpms
```

- For Red Hat Enterprise Linux 8

```
# subscription-manager repos --enable=rhel-8-for-x86_64-appstream-rpms
```

3. Install the guest agent and dependencies:

- For Red Hat Enterprise Linux 6 or 7, install the ovirt guest agent:

```
# yum install ovirt-guest-agent-common
```

- For Red Hat Enterprise Linux 8 and 9, install the qemu guest agent:

```
# yum install qemu-guest-agent
```

4. Start and enable the **ovirt-guest-agent** service:

- For Red Hat Enterprise Linux 6

```
# service ovirt-guest-agent start  
# chkconfig ovirt-guest-agent on
```

- For Red Hat Enterprise Linux 7

```
# systemctl start ovirt-guest-agent  
# systemctl enable ovirt-guest-agent
```

5. Start and enable the **qemu-guest-agent** service:

- For Red Hat Enterprise Linux 6

```
# service qemu-ga start  
# chkconfig qemu-ga on
```

- For Red Hat Enterprise Linux 7, 8 or 9

```
# systemctl start qemu-guest-agent  
# systemctl enable qemu-guest-agent
```

The guest agent now passes usage information to the Red Hat Virtualization Manager. You can configure the oVirt guest agent in the **/etc/ovirt-guest-agent.conf** file.

CHAPTER 3. INSTALLING WINDOWS VIRTUAL MACHINES

Installing a Windows virtual machine involves the following key steps:

1. Create a blank virtual machine on which to install an operating system.
2. Add a virtual disk for storage.
3. Add a network interface to connect the virtual machine to the network.
4. Attach the Windows guest tools CD to the virtual machine so that VirtIO-optimized device drivers can be installed during the operating system installation.
5. Install a Windows operating system on the virtual machine. See your operating system's documentation for instructions.
6. During the installation, install guest agents and drivers for additional virtual machine functionality.

When all of these steps are complete, the new virtual machine is functional and ready to perform tasks.

3.1. CREATING A VIRTUAL MACHINE

When creating a new virtual machine, you specify its settings. You can edit some of these settings later, including the chipset and BIOS type. For more information, see [UEFI and the Q35 chipset](#) in the *Administration Guide*. .Prerequisites



NOTE

Before you can use this virtual machine, you must:

- Install an operating system
- Install a VirtIO-optimized disk and network drivers

Procedure

1. You can change the default virtual machine name length with the **engine-config** tool. Run the following command on the Manager machine:

```
# engine-config --set MaxVmNameLength=integer
```

2. Click **Compute** → **Virtual Machines**.
3. Click **New**. This opens the **New Virtual Machine** window.
4. Select an **Operating System** from the drop-down list.
5. Enter a **Name** for the virtual machine.
6. Add storage to the virtual machine: under **Instance Images**, click **Attach** or **Create** to select or create a virtual disk .
 - Click **Attach** and select an existing virtual disk.
 - or

- Click **Create** and enter a **Size(GB)** and **Alias** for a new virtual disk. You can accept the default settings for all other fields, or change them if required. See [Explanation of settings in the New Virtual Disk and Edit Virtual Disk windows](#) for more details on the fields for all disk types.
7. Connect the virtual machine to the network. Add a network interface by selecting a vNIC profile from the **nic1** drop-down list at the bottom of the **General** tab.
 8. Specify the virtual machine's **Memory Size** on the **System** tab.
 9. In the **Boot Options** tab, choose the **First Device** that the virtual machine will use to boot.
 10. You can accept the default settings for all other fields, or change them if required. For more details on all fields in the **New Virtual Machine** window, see [Explanation of settings in the New Virtual Machine and Edit Virtual Machine Windows](#).
 11. Click **OK**.

The new virtual machine is created and displays in the list of virtual machines with a status of **Down**.

3.2. STARTING THE VIRTUAL MACHINE USING RUN ONCE

3.2.1. Installing Windows on VirtIO-optimized hardware

Install VirtIO-optimized disk and network device drivers during your Windows installation by attaching the **virtio-win_version.iso** file to your virtual machine. These drivers provide a performance improvement over emulated device drivers.

Use the **Run Once** option to attach the **virtio-win_version.iso** file in a one-off boot different from the **Boot Options** defined in the **New Virtual Machine** window.

Prerequisites

The following items were added to the virtual machine:

- A **Red Hat VirtIO** network interface
- A disk that uses the **VirtIO** interface

You can upload **virtio-win_version.iso** to a data storage domain.



NOTE

Red Hat recommends uploading ISO images to the data domain with the Administration Portal or with the REST API. For more information, see [Uploading Images to a Data Storage Domain](#) in the *Administration Guide*.

If necessary, you can upload the **virtio-win** ISO file to an ISO storage domain that is hosted on the Manager. The ISO storage domain type is deprecated. For more information, see [Uploading images to an ISO domain](#) in the *Administration Guide*.

Procedure

To install the virtio-win drivers when installing Windows, complete the following steps:

1. Click **Compute** → **Virtual Machines** and select a virtual machine.

2. Click **Run** → **Run Once**.
3. Expand the **Boot Options** menu.
4. Select the **Attach CD** check box, and select a Windows ISO from the drop-down list.
5. Select the **Attach Windows guest tools CD** check box.
6. Move **CD-ROM** to the top of the **Boot Sequence** field.
7. Configure other **Run Once** options as required. See [Virtual Machine Run Once settings explained](#) for more details.
8. Click **OK**. The status of the virtual machine changes to **Up**, and the operating system installation begins.
Open a console to the virtual machine if one does not open automatically during the Windows installation.
9. When prompted to select a drive onto which you want to install Windows, click **Load driver** and **OK**.
10. Under **Select the driver to install**, select the appropriate driver for the version of Windows. For example, for Windows Server 2019, select **Red Hat VirtIO SCSI controller (E:\amd64\2k19\viostor.inf)**
11. Click Next.

The rest of the installation proceeds as normal.

3.2.2. Opening a console to a virtual machine

Use Remote Viewer to connect to a virtual machine.



NOTE

To allow other users to connect to the VM, make sure you shutdown and restart the virtual machine when you are finished using the console. Alternatively, the administrator can **Disable strict user checking** to eliminate the need for reboot between users. See [Virtual Machine Console Settings Explained](#) for more information.

Procedure

1. Install Remote Viewer if it is not already installed. See [Installing Console Components](#).
2. Click **Compute** → **Virtual Machines** and select a virtual machine.
3. Click **Console**. By default, the browser prompts you to download a file named **console.vv**. When you click to open the file, a console window opens for the virtual machine. You can configure your browser to automatically open these files, such that clicking **Console** simply opens the console.



NOTE

console.vv expires after 120 seconds. If more than 120 seconds elapse between the time the file is downloaded and the time that you open the file, click **Console** again.

Additional resources

- [Automatically connecting to a Virtual Machine](#)
- [Configuring Console Options](#)

3.3. INSTALLING GUEST AGENTS AND DRIVERS

3.3.1. Red Hat Virtualization Guest agents, tools, and drivers

The Red Hat Virtualization guest agents, tools, and drivers provide additional functionality for virtual machines, such as gracefully shutting down or rebooting virtual machines from the VM Portal and Administration Portal. The tools and agents also provide information for virtual machines, including:

- Resource usage
- IP addresses

The guest agents, tools and drivers are distributed as an ISO file that you can attach to virtual machines. This ISO file is packaged as an RPM file that you can install and upgrade from the Manager machine.

You need to install the guest agents and drivers on a virtual machine to enable this functionality for that machine.

Table 3.1. Red Hat Virtualization Guest drivers

Driver	Description	Works on
virtio-net	Paravirtualized network driver provides enhanced performance over emulated devices like rtl.	Server and Desktop.
virtio-block	Paravirtualized HDD driver offers increased I/O performance over emulated devices like IDE by optimizing the coordination and communication between the virtual machine and the hypervisor. The driver complements the software implementation of the virtio-device used by the host to play the role of a hardware device.	Server and Desktop.
virtio-scsi	Paravirtualized iSCSI HDD driver offers similar functionality to the virtio-block device, with some additional enhancements. In particular, this driver supports adding hundreds of devices, and names devices using the standard SCSI device naming scheme.	Server and Desktop.

Driver	Description	Works on
virtio-serial	Virtio-serial provides support for multiple serial ports. The improved performance is used for fast communication between the virtual machine and the host that avoids network complications. This fast communication is required for the guest agents and for other features such as clipboard copy-paste between the virtual machine and the host and logging.	Server and Desktop.
virtio-balloon	Virtio-balloon is used to control the amount of memory a virtual machine actually accesses. It offers improved memory overcommitment.	Server and Desktop.
qxl	A paravirtualized display driver reduces CPU usage on the host and provides better performance through reduced network bandwidth on most workloads.	Server and Desktop.

Table 3.2. Red Hat Virtualization Guest agents and tools

Guest agent/tool	Description	Works on
qemu-guest-agent	Used instead of ovirt-guest-agent-common on Red Hat Enterprise Linux 8 virtual machines. It is installed and enabled by default.	Server and Desktop.

Guest agent/tool	Description	Works on
spice-agent	<p>The SPICE agent supports multiple monitors and is responsible for client-mouse-mode support to provide a better user experience and improved responsiveness than the QEMU emulation. Cursor capture is not needed in client-mouse-mode. The SPICE agent reduces bandwidth usage when used over a wide area network by reducing the display level, including color depth, disabling wallpaper, font smoothing, and animation. The SPICE agent enables clipboard support allowing cut and paste operations for both text and images between client and virtual machine, and automatic guest display setting according to client-side settings. On Windows-based virtual machines, the SPICE agent consists of vdservice and vdagent.</p>	Server and Desktop.

3.3.2. Installing the guest agents, tools, and drivers on Windows

Procedure

To install the guest agents, tools, and drivers on a Windows virtual machine, complete the following steps:

1. On the Manager machine, install the **virtio-win** package:

```
# dnf install virtio-win*
```

After you install the package, the ISO file is located in **/usr/share/virtio-win/virtio-win_ *version*.iso** on the Manager machine.

2. Upload **virtio-win_ *version*.iso** to a data storage domain. See [Uploading Images to a Data Storage Domain](#) in the *Administration Guide* for details.
3. In the Administration or VM Portal, if the virtual machine is running, use the **Change CD** button to attach the **virtio-win_ *version*.iso** file to each of your virtual machines. If the virtual machine is powered off, click the **Run Once** button and attach the ISO as a CD.
4. Log in to the virtual machine.
5. Select the CD Drive containing the **virtio-win_ *version*.iso** file. You can complete the installation with either the GUI or the command line.
6. Run the installer.

To install with the GUI, complete the following steps

- a. Double-click **virtio-win-guest-tools.exe**.
- b. Click **Next** at the welcome screen.
- c. Follow the prompts in the installation wizard.
- d. When installation is complete, select **Yes, I want to restart my computer now** and click **Finish** to apply the changes.

To install silently with the command line, complete the following steps

- a. Open a command prompt with Administrator privileges.
- b. Enter the **msiexec** command:

```
D:\ msiexec /i "PATH_TO_MSI" /qn [/!*"PATH_TO_LOG"][/norestart]
ADDLOCAL=ALL
```

Other possible values for ADDLOCAL are listed below.

For example, to run the installation when **virtio-win-gt-x64.msi** is on the **D:** drive, without saving the log, and then immediately restart the virtual machine, enter the following command:

```
D:\ msiexec /i "virtio-win-gt-x64.msi" /qn ADDLOCAL=ALL
```

After installation completes, the guest agents and drivers pass usage information to the Red Hat Virtualization Manager and enable you to access USB devices and other functionality.

3.3.3. Values for ADDLOCAL to customize virtio-win command-line installation

When installing **virtio-win-gt-x64.msi** or **virtio-win-gt-x32.msi** with the command line, you can install any one driver, or any combination of drivers.

You can also install specific agents, but you must also install each agent's corresponding drivers.

The **ADDLOCAL** parameter of the **msiexec** command enables you to specify which drivers or agents to install. **ADDLOCAL=ALL** installs all drivers and agents. Other values are listed in the following tables.

Table 3.3. Possible values for ADDLOCAL to install drivers

Value for ADDLOCAL	Driver Name	Description
FE_network_driver	virtio-net	Paravirtualized network driver provides enhanced performance over emulated devices like rtl.
FE_balloon_driver	virtio-balloon	Controls the amount of memory a virtual machine actually accesses. It offers improved memory overcommitment.

Value for ADDLOCAL	Driver Name	Description
FE_pvpanic_driver	pvpanic	QEMU pvpanic device driver.
FE_qemufwcfg_driver	qemufwcfg	QEMU FWCfg device driver.
FE_qemupcserial_driver	qemupcserial	QEMU PCI serial device driver.
FE_spice_driver	Spice Driver	A paravirtualized display driver reduces CPU usage on the host and provides better performance through reduced network bandwidth on most workloads.
FE_vioinput_driver	vioinput	VirtIO Input Driver.
FE_viorng_driver	viorng	VirtIO RNG device driver.
FE_vioscsi_driver	vioscsi	VirtIO SCSI pass-through controller.
FE_vioserial_driver	vioserial	VirtIO Serial device driver.
FE_viostor_driver	viostor	VirtIO Block driver.

Table 3.4. Possible values for ADDLOCAL to install agents and required corresponding drivers

Agent	Description	Corresponding driver(s)	Value for ADDLOCAL
Spice Agent	Supports multiple monitors, responsible for client-mouse-mode support, reduces bandwidth usage, enables clipboard support between client and virtual machine, provide a better user experience and improved responsiveness.	vioserial and Spice driver	FE_spice_Agent,FE_vioserial_driver,FE_spice_driver

Examples

The following command installs only the VirtIO SCSI pass-through controller, the VirtIO Serial device driver, and the VirtIO Block driver:

```
D:\ msiexec /i "virtio-win-gt-x64.msi" /qn  
ADDLOCAL=`FE_vioscsi_driver,FE_vioserial_driver,FE_viostor_driver
```

The following command installs only the Spice Agent and its required corresponding drivers:

```
D:\ msiexec /i "virtio-win-gt-x64.msi" /qn ADDLOCAL =  
FE_spice_Agent,FE_vioserial_driver,FE_spice_driver
```

Additional resources

- [Updating Win Guest Drivers with Windows Updates](#)
- [Updating the Guest Agents and Drivers on Windows](#)

The Microsoft Developer website:

- [Windows Installer](#)
- [Command-Line Options](#) for the Windows installer
- [Property Reference](#) for the Windows installer

CHAPTER 4. ADDITIONAL CONFIGURATION

4.1. CONFIGURING OPERATING SYSTEMS WITH OSINFO

Red Hat Virtualization stores operating system configurations for virtual machines in `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties`. This file contains default values such as `os.other.devices.display.protocols.value = spice/qxl,vnc/vga,vnc/qxl`.

There are only a limited number of scenarios in which you would change these values:

- Adding an operating system that does not appear in the list of supported guest operating systems
- Adding a product key (for example, `os.windows_10x64.productKey.value =`)
- Configuring the `sysprep` path for a Windows virtual machine (for example, `os.windows_10x64.sysprepPath.value = ${ENGINE_USR}/conf/sysprep/sysprep.w10x64`)



IMPORTANT

Do not edit the actual `00-defaults.properties` file. Changes will be overwritten if you upgrade or restore the Manager.

Do not change values that come directly from the operating system or the Manager, such as maximum memory size.

To change the operating system configurations, create an override file in `/etc/ovirt-engine/osinfo.conf.d/`. The file name must begin with a value greater than `00`, so that the file appears after `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties`, and ends with the extension, `.properties`.

For example, `10-productkeys.properties` overrides the default file, `00-defaults.properties`. The last file in the file list has precedence over earlier files.

For virtual machines with **Windows Server 2016**, on each RHV host, modify the `/etc/modprobe.d/kvm.conf` configuration file by adding the line: `options kvm ignore_msrs=1`


4.2. CONFIGURING SINGLE SIGN-ON FOR VIRTUAL MACHINES

Configuring single sign-on, also known as password delegation, allows you to automatically log in to a virtual machine using the credentials you use to log in to the VM Portal. Single sign-on can be used on both Red Hat Enterprise Linux and Windows virtual machines.



NOTE

Single sign-on is not supported for virtual machines running Red Hat Enterprise Linux 8.0.



IMPORTANT

If single sign-on to the VM Portal is enabled, single sign-on to virtual machines will not be possible. With single sign-on to the VM Portal enabled, the VM Portal does not need to accept a password, thus the password cannot be delegated to sign in to virtual machines.

4.2.1. Configuring Single Sign-On for Red Hat Enterprise Linux Virtual Machines Using IPA (IdM)

To configure single sign-on for Red Hat Enterprise Linux virtual machines using GNOME and KDE graphical desktop environments and IPA (IdM) servers, you must install the **ovirt-guest-agent** package on the virtual machine and install the packages associated with your window manager.



IMPORTANT

The following procedure assumes that you have a working IPA configuration and that the IPA domain is already joined to the Manager. You must also ensure that the clocks on the Manager, the virtual machine and the system on which IPA (IdM) is hosted are synchronized using NTP.



NOTE

Single sign-on with IPA (IdM) is deprecated for virtual machines running Red Hat Enterprise Linux version 7 or earlier and unsupported for virtual machines running Red Hat Enterprise Linux 8 or Windows operating systems.

Configuring Single Sign-On for Red Hat Enterprise Linux Virtual Machines

1. Log in to the Red Hat Enterprise Linux virtual machine.
2. Enable the repository:

- For Red Hat Enterprise Linux 6:

```
# subscription-manager repos --enable=rhel-6-server-rhv-4-agent-rpms
```

- For Red Hat Enterprise Linux 7:

```
# subscription-manager repos --enable=rhel-7-server-rh-common-rpms
```

3. Download and install the guest agent, single sign-on, and IPA packages:

```
# yum install ovirt-guest-agent-common ovirt-guest-agent-pam-module ovirt-guest-agent-gdm-plugin ipa-client
```

4. Run the following command and follow the prompts to configure **ipa-client** and join the virtual machine to the domain:

```
# ipa-client-install --permit --mkhomedir
```




NOTE

In environments that use DNS obfuscation, this command should be:

```
# ipa-client-install --domain=FQDN --server=FQDN
```

5. For Red Hat Enterprise Linux 7.2 and later:


```
# authconfig --enablenis --update
```



NOTE

Red Hat Enterprise Linux 7.2 has a new version of the System Security Services Daemon (SSSD), which introduces configuration that is incompatible with the Red Hat Virtualization Manager guest agent single sign-on implementation. This command ensures that single sign-on works.

6. Fetch the details of an IPA user:

```
# getent passwd ipa-user
```

7. Record the IPA user's UID and GID:

```
ipa-user:*:936600010:936600001::/home/ipa-user:/bin/sh
```

8. Create a home directory for the IPA user:

```
# mkdir /home/ipa-user
```

9. Assign ownership of the directory to the IPA user:

```
# chown 936600010:936600001 /home/ipa-user
```

Log in to the VM Portal using the user name and password of a user configured to use single sign-on and connect to the console of the virtual machine. You will be logged in automatically.

4.2.2. Configuring single sign-on for Windows virtual machines

To configure single sign-on for Windows virtual machines, the Windows guest agent must be installed on the guest virtual machine. The `virtio-win` ISO image provides this agent. If the **`virtio-win_version.iso`** image is not available in your storage domain, contact your system administrator.

Procedure

1. Select the Windows virtual machine. Ensure the machine is powered up.
2. On the virtual machine, locate the CD drive and open the CD.
3. Launch **`virtio-win-guest-tools`**.
4. Click **Options**
5. Select **Install oVirt Guest Agent**
6. Click **OK**.
7. Click **Install**.
8. When the installation completes, you are prompted to restart the machine to apply the changes.

Log in to the VM Portal using the user name and password of a user configured to use single sign-on and connect to the console of the virtual machine. You will be logged in automatically.

4.2.3. Disabling Single Sign-on for Virtual Machines

The following procedure explains how to disable single sign-on for a virtual machine.

Disabling Single Sign-On for Virtual Machines

1. Select a virtual machine and click **Edit**.
2. Click the **Console** tab.
3. Select the **Disable Single Sign On** check box.
4. Click **OK**.

4.3. CONFIGURING USB DEVICES

A virtual machine connected with the SPICE protocol can be configured to connect directly to USB devices.

The USB device will only be redirected if the virtual machine is active, in focus and is run from the VM Portal. USB redirection can be manually enabled each time a device is plugged in or set to automatically redirect to active virtual machines in the Console Options window.



IMPORTANT

Note the distinction between the client machine and guest machine. The client is the hardware from which you access a guest. The guest is the virtual desktop or virtual server which is accessed through the VM Portal or Administration Portal.

USB redirection **Enabled** mode allows KVM/SPICE USB redirection for Linux and Windows virtual machines. Virtual (guest) machines require no guest-installed agents or drivers for native USB. On Red Hat Enterprise Linux clients, all packages required for USB redirection are provided by the **virt-viewer** package. On Windows clients, you must also install the **usbdk** package. Enabled USB mode is supported on the following clients and guests:




NOTE

If you have a 64-bit architecture PC, you must use the 64-bit version of Internet Explorer to install the 64-bit version of the USB driver. The USB redirection will not work if you install the 32-bit version on a 64-bit architecture. As long as you initially install the correct USB type, you can access USB redirection from both 32- and 64-bit browsers.

4.3.1. Using USB Devices on a Windows Client

The **usbdk** driver must be installed on the Windows client for the USB device to be redirected to the guest. Ensure the version of **usbdk** matches the architecture of the client machine. For example, the 64-bit version of **usbdk** must be installed on 64-bit Windows machines.



NOTE


USB redirection is only supported when you open the virtual machine from the VM Portal.

Procedure

1. When the **usbdk** driver is installed, click **Compute → Virtual Machines** and select a virtual machine that is configured to use the SPICE protocol.
2. Click the **Console** tab.
3. Select the **USB enabled** checkbox and click **OK**.
4. Click **Console → Console Options**.
5. Select the **Enable USB Auto-Share** check box and click **OK**.
6. Start the virtual machine from the VM Portal and click **Console** to connect to that virtual machine.
7. Plug your USB device into the client machine to make it appear automatically on the guest machine.

4.3.2. Using USB Devices on a Red Hat Enterprise Linux Client

The **usbredir** package enables USB redirection from Red Hat Enterprise Linux clients to virtual machines. **usbredir** is a dependency of the **virt-viewer** package, and is automatically installed together with that package.

**NOTE**

USB redirection is only supported when you open the virtual machine from the VM Portal.

Procedure

1. Click **Compute → Virtual Machines**.
2. Select a virtual machine that has been configured to use the SPICE protocol and click **Edit**. This opens the **Edit Virtual Machine** window.
3. Click the **Console** tab.
4. Select the **USB enabled** checkbox and click **OK**.
5. Click **Console → Console Options**.
6. Select the **Enable USB Auto-Share** check box and click **OK**.
7. Start the virtual machine from the VM Portal and click **Console** to connect to that virtual machine.
8. Plug your USB device into the client machine to make it appear automatically on the guest machine.

4.4. CONFIGURING MULTIPLE MONITORS**4.4.1. Configuring Multiple Displays for Red Hat Enterprise Linux Virtual Machines**

A maximum of four displays can be configured for a single Red Hat Enterprise Linux virtual machine when connecting to the virtual machine using the SPICE protocol.

1. Start a SPICE session with the virtual machine.
2. Open the **View** drop-down menu at the top of the SPICE client window.
3. Open the **Display** menu.
4. Click the name of a display to enable or disable that display.




NOTE

By default, **Display 1** is the only display that is enabled on starting a SPICE session with a virtual machine. If no other displays are enabled, disabling this display will close the session.

4.4.2. Configuring Multiple Displays for Windows Virtual Machines

A maximum of four displays can be configured for a single Windows virtual machine when connecting to the virtual machine using the SPICE protocol.

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. With the virtual machine in a powered-down state, click **Edit**.
3. Click the **Console** tab.
4. Select the number of displays from the **Monitors** drop-down list.



NOTE

This setting controls the maximum number of displays that can be enabled for the virtual machine. While the virtual machine is running, additional displays can be enabled up to this number.

5. Click **OK**.
6. Start a SPICE session with the virtual machine.
7. Open the **View** drop-down menu at the top of the SPICE client window.
8. Open the **Display** menu.
9. Click the name of a display to enable or disable that display.



NOTE

By default, **Display 1** is the only display that is enabled on starting a SPICE session with a virtual machine. If no other displays are enabled, disabling this display will close the session.

4.5. CONFIGURING CONSOLE OPTIONS

4.5.1. Console Options

Connection protocols are the underlying technology used to provide graphical consoles for virtual machines and allow users to work with virtual machines in a similar way as they would with physical machines. Red Hat Virtualization currently supports the following connection protocols:

SPICE

Simple Protocol for Independent Computing Environments (SPICE) is the recommended connection protocol for both Linux virtual machines and Windows virtual machines. To open a console to a virtual machine using SPICE, use Remote Viewer.

VNC

Virtual Network Computing (VNC) can be used to open consoles to both Linux virtual machines and Windows virtual machines. To open a console to a virtual machine using VNC, use Remote Viewer or a VNC client.

RDP

Remote Desktop Protocol (RDP) can only be used to open consoles to Windows virtual machines, and is only available when you access a virtual machines from a Windows machine on which Remote Desktop has been installed. Before you can connect to a Windows virtual machine using RDP, you must set up remote sharing on the virtual machine and configure the firewall to allow remote desktop connections.



NOTE

SPICE is not supported on virtual machines running Windows 8 or Windows 8.1. If a virtual machine running one of these operating systems is configured to use the SPICE protocol, it detects the absence of the required SPICE drivers and runs in VGA compatibility mode.

4.5.2. Accessing Console Options

You can configure several options for opening graphical consoles for virtual machines in the Administration Portal.

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Console** → **Console Options**.



NOTE

You can configure the connection protocols and video type in the **Console** tab of the **Edit Virtual Machine** window in the Administration Portal. Additional options specific to each of the connection protocols, such as the keyboard layout when using the VNC connection protocol, can be configured. See [Virtual Machine Console settings explained](#) for more information.

4.5.3. SPICE Console Options

When the SPICE connection protocol is selected, the following options are available in the **Console Options** window.

SPICE Options

- **Map control-alt-del shortcut to ctrl+alt+end** Select this check box to map the **Ctrl + Alt + Del** key combination to **Ctrl + Alt + End** inside the virtual machine.
- **Enable USB Auto-Share:** Select this check box to automatically redirect USB devices to the virtual machine. If this option is not selected, USB devices will connect to the client machine instead of the guest virtual machine. To use the USB device on the guest machine, manually enable it in the SPICE client menu.
- **Open in Full Screen** Select this check box for the virtual machine console to automatically open in full screen when you connect to the virtual machine. Press **SHIFT + F11** to toggle full screen mode on or off.
- **Enable SPICE Proxy:** Select this check box to enable the SPICE proxy.

4.5.4. VNC Console Options

When the VNC connection protocol is selected, the following options are available in the **Console Options** window.

Console Invocation

- **Native Client:** When you connect to the console of the virtual machine, a file download dialog provides you with a file that opens a console to the virtual machine via Remote Viewer.
- **noVNC:** When you connect to the console of the virtual machine, a browser tab is opened that acts as the console.

VNC Options

- **Map control-alt-delete shortcut to ctrl+alt+end** Select this check box to map the **Ctrl + Alt + Del** key combination to **Ctrl + Alt + End** inside the virtual machine.

4.5.5. RDP Console Options

When the RDP connection protocol is selected, the following options are available in the **Console Options** window.

Console Invocation

- **Auto:** The Manager automatically selects the method for invoking the console.
- **Native client:** When you connect to the console of the virtual machine, a file download dialog provides you with a file that opens a console to the virtual machine via Remote Desktop.

RDP Options

- **Use Local Drives:** Select this check box to make the drives on the client machine accessible on the guest virtual machine.

4.5.6. Remote Viewer Options

4.5.6.1. Remote Viewer Options

When you specify the **Native client** console invocation option, you will connect to virtual machines using Remote Viewer. The Remote Viewer window provides a number of options for interacting with the virtual machine to which it is connected.

Table 4.1. Remote Viewer Options

Option	Hotkey
File	<ul style="list-style-type: none"> ● Screenshot: Takes a screen capture of the active window and saves it in a location of your specification. ● USB device selection: If USB redirection has been enabled on your virtual machine, the USB device plugged into your client machine can be accessed from this menu. ● Quit: Closes the console. The hot key for this option is Shift + Ctrl + Q.
View	<ul style="list-style-type: none"> ● Full screen: Toggles full screen mode on or off. When enabled, full screen mode expands the virtual machine to fill the entire screen. When disabled, the virtual machine is displayed as a window. The hot key for enabling or disabling full screen is SHIFT + F11. ● Zoom: Zooms in and out of the console window. Ctrl + + zooms in, Ctrl + - zooms out, and Ctrl + 0 returns the screen to its original size. ● Automatically resize: Select to enable the guest resolution to automatically scale according to the size of the console window. ● Displays: Allows users to enable and disable displays for the guest virtual machine.

Option	Hotkey
Send key	<ul style="list-style-type: none"> ● Ctrl + Alt + Del: On a Red Hat Enterprise Linux virtual machine, it displays a dialog with options to suspend, shut down or restart the virtual machine. On a Windows virtual machine, it displays the task manager or Windows Security dialog. ● Ctrl + Alt + Backspace: On a Red Hat Enterprise Linux virtual machine, it restarts the X sever. On a Windows virtual machine, it does nothing. ● Ctrl + Alt + F1 ● Ctrl + Alt + F2 ● Ctrl + Alt + F3 ● Ctrl + Alt + F4 ● Ctrl + Alt + F5 ● Ctrl + Alt + F6 ● Ctrl + Alt + F7 ● Ctrl + Alt + F8 ● Ctrl + Alt + F9 ● Ctrl + Alt + F10 ● Ctrl + Alt + F11 ● Ctrl + Alt + F12 ● Printscreen: Passes the Printscreen keyboard option to the virtual machine.
Help	The About entry displays the version details of Virtual Machine Viewer that you are using.
Release Cursor from Virtual Machine	SHIFT + F12

4.5.6.2. Remote Viewer Hotkeys

You can access the hotkeys for a virtual machine in both full screen mode and windowed mode. If you are using full screen mode, you can display the menu containing the button for hotkeys by moving the mouse pointer to the middle of the top of the screen. If you are using windowed mode, you can access the hotkeys via the **Send key** menu on the virtual machine window title bar.

**NOTE**

If **vdagent** is not running on the client machine, the mouse can become captured in a virtual machine window if it is used inside a virtual machine and the virtual machine is not in full screen. To unlock the mouse, press **Shift + F12**.

4.5.6.3. Manually Associating console.vv Files with Remote Viewer

If you are prompted to download a **console.vv** file when attempting to open a console to a virtual machine using the native client console option, and Remote Viewer is already installed, then you can manually associate **console.vv** files with Remote Viewer so that Remote Viewer can automatically use those files to open consoles.

Manually Associating console.vv Files with Remote Viewer

1. Start the virtual machine.
2. Open the **Console Options** window:
 - In the Administration Portal, click **Console → Console Options**.
 - In the VM Portal, click the virtual machine name and click the pencil icon beside **Console**.
3. Change the console invocation method to **Native client** and click **OK**.
4. Attempt to open a console to the virtual machine, then click **Save** when prompted to open or save the **console.vv** file.
5. Click the location on your local machine where you saved the file.
6. Double-click the **console.vv** file and select **Select a program from a list of installed programs** when prompted.
7. In the **Open with** window, select **Always use the selected program to open this kind of file** and click the **Browse** button.
8. Click the **C:\Users_[user name]_\AppData\Local\virt-viewer\bin** directory and select **remote-viewer.exe**.
9. Click **Open** and then click **OK**.

When you use the native client console invocation option to open a console to a virtual machine, Remote Viewer will automatically use the **console.vv** file that the Red Hat Virtualization Manager provides to open a console to that virtual machine without prompting you to select the application to use.

4.6. CONFIGURING A WATCHDOG

4.6.1. Adding a Watchdog Card to a Virtual Machine

You can add a watchdog card to a virtual machine to monitor the operating system's responsiveness.

Procedure

1. Click **Compute → Virtual Machines** and select a virtual machine.
2. Click **Edit**.

3. Click the **High Availability** tab.
4. Select the watchdog model to use from the **Watchdog Model** drop-down list.
5. Select an action from the **Watchdog Action** drop-down list. This is the action that the virtual machine takes when the watchdog is triggered.
6. Click **OK**.

4.6.2. Installing a Watchdog

To activate a watchdog card attached to a virtual machine, you must install the **watchdog** package on that virtual machine and start the **watchdog** service.

Installing Watchdogs

1. Log in to the virtual machine on which the watchdog card is attached.
2. Install the **watchdog** package and dependencies:

```
# yum install watchdog
```

3. Edit the **/etc/watchdog.conf** file and uncomment the following line:

```
watchdog-device = /dev/watchdog
```

4. Save the changes.
5. Start the **watchdog** service and ensure this service starts on boot:

- Red Hat Enterprise Linux 6:

```
# service watchdog start  
# chkconfig watchdog on
```

- Red Hat Enterprise Linux 7:

```
# systemctl start watchdog.service  
# systemctl enable watchdog.service
```

4.6.3. Confirming Watchdog Functionality

Confirm that a watchdog card has been attached to a virtual machine and that the **watchdog** service is active.



WARNING

This procedure is provided for testing the functionality of watchdogs only and must not be run on production machines.

Confirming Watchdog Functionality

1. Log in to the virtual machine on which the watchdog card is attached.
2. Confirm that the watchdog card has been identified by the virtual machine:

```
# lspci | grep watchdog -i
```

3. Run one of the following commands to confirm that the watchdog is active:

- Trigger a kernel panic:

```
# echo c > /proc/sysrq-trigger
```

- Terminate the **watchdog** service:

```
# kill -9 pgrep watchdog
```

The watchdog timer can no longer be reset, so the watchdog counter reaches zero after a short period of time. When the watchdog counter reaches zero, the action specified in the **Watchdog Action** drop-down menu for that virtual machine is performed.

4.6.4. Parameters for Watchdogs in `watchdog.conf`

The following is a list of options for configuring the **watchdog** service available in the `/etc/watchdog.conf` file. To configure an option, you must uncomment that option and restart the **watchdog** service after saving the changes.



NOTE

For a more detailed explanation of options for configuring the **watchdog** service and using the **watchdog** command, see the **watchdog** man page.

Table 4.2. `watchdog.conf` variables

Variable name	Default Value	Remarks
ping	N/A	An IP address that the watchdog attempts to ping to verify whether that address is reachable. You can specify multiple IP addresses by adding additional ping lines.
interface	N/A	A network interface that the watchdog will monitor to verify the presence of network traffic. You can specify multiple network interfaces by adding additional interface lines.

Variable name	Default Value	Remarks
file	/var/log/messages	A file on the local system that the watchdog will monitor for changes. You can specify multiple files by adding additional file lines.
change	1407	The number of watchdog intervals after which the watchdog checks for changes to files. A change line must be specified on the line directly after each file line, and applies to the file line directly above that change line.
max-load-1	24	The maximum average load that the virtual machine can sustain over a one-minute period. If this average is exceeded, then the watchdog is triggered. A value of 0 disables this feature.
max-load-5	18	The maximum average load that the virtual machine can sustain over a five-minute period. If this average is exceeded, then the watchdog is triggered. A value of 0 disables this feature. By default, the value of this variable is set to a value approximately three quarters that of max-load-1 .
max-load-15	12	The maximum average load that the virtual machine can sustain over a fifteen-minute period. If this average is exceeded, then the watchdog is triggered. A value of 0 disables this feature. By default, the value of this variable is set to a value approximately one half that of max-load-1 .
min-memory	1	The minimum amount of virtual memory that must remain free on the virtual machine. This value is measured in pages. A value of 0 disables this feature.

Variable name	Default Value	Remarks
repair-binary	/usr/sbin/repair	The path and file name of a binary file on the local system that will be run when the watchdog is triggered. If the specified file resolves the issues preventing the watchdog from resetting the watchdog counter, then the watchdog action is not triggered.
test-binary	N/A	The path and file name of a binary file on the local system that the watchdog will attempt to run during each interval. A test binary allows you to specify a file for running user-defined tests.
test-timeout	N/A	The time limit, in seconds, for which user-defined tests can run. A value of 0 allows user-defined tests to continue for an unlimited duration.
temperature-device	N/A	The path to and name of a device for checking the temperature of the machine on which the watchdog service is running.
max-temperature	120	The maximum allowed temperature for the machine on which the watchdog service is running. The machine will be halted if this temperature is reached. Unit conversion is not taken into account, so you must specify a value that matches the watchdog card being used.
admin	root	The email address to which email notifications are sent.


Variable name	Default Value	Remarks
interval	10	The interval, in seconds, between updates to the watchdog device. The watchdog device expects an update at least once every minute, and if there are no updates over a one-minute period, then the watchdog is triggered. This one-minute period is hard-coded into the drivers for the watchdog device, and cannot be configured.
logtick	1	When verbose logging is enabled for the watchdog service, the watchdog service periodically writes log messages to the local system. The logtick value represents the number of watchdog intervals after which a message is written.
realtime	yes	Specifies whether the watchdog is locked in memory. A value of yes locks the watchdog in memory so that it is not swapped out of memory, while a value of no allows the watchdog to be swapped out of memory. If the watchdog is swapped out of memory and is not swapped back in before the watchdog counter reaches zero, then the watchdog is triggered.
priority	1	The schedule priority when the value of realtime is set to yes .
pidfile	/var/run/syslogd.pid	The path and file name of a PID file that the watchdog monitors to see if the corresponding process is still active. If the corresponding process is not active, then the watchdog is triggered.

4.7. CONFIGURING VIRTUAL NUMA

In the Administration Portal, you can configure virtual NUMA nodes on a virtual machine and pin them to physical NUMA nodes on one or more hosts. The host's default policy is to schedule and run virtual machines on any available resources on the host. As a result, the resources backing a large virtual

machine that cannot fit within a single host socket could be spread out across multiple NUMA nodes. Over time these resources may be moved around, leading to poor and unpredictable performance. Configure and pin virtual NUMA nodes to avoid this outcome and improve performance.

Configuring virtual NUMA requires a NUMA-enabled host. To confirm whether NUMA is enabled on a host, log in to the host and run **numactl --hardware**. The output of this command should show at least two NUMA nodes. You can also view the host's NUMA topology in the Administration Portal by selecting the host from the **Hosts** tab and clicking **NUMA Support**. This button is only available when the selected host has at least two NUMA nodes.



NOTE

If you define **NUMA Pinning**, the default migration mode is **Allow manual migration only** by default.

Configuring Virtual NUMA

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click **Show Advanced Options**.
4. Click the **Host** tab.
5. Select the **Specific Host(s)** radio button and select the host(s) from the list. The selected host(s) must have at least two NUMA nodes.
6. Click **NUMA Pinning**.
7. In the **NUMA Topology** window, click and drag virtual NUMA nodes from the box on the right to host NUMA nodes on the left as required, and click **OK**.
8. Select **Strict**, **Preferred**, or **Interleave** from the **Tune Mode** drop-down list in each NUMA node. If the selected mode is **Preferred**, the **NUMA Node Count** must be set to **1**.
9. You can also set the NUMA pinning policy automatically by selecting **Resize and Pin NUMA** from the CPU Pinning Policy drop-down list under the **CPU Allocation** settings in the **Resource Allocation** tab:
 - **None** - Runs without any CPU pinning.
 - **Manual** - Runs a manually specified virtual CPU on a specific physical CPU and a specific host. Available only when the virtual machine is pinned to a Host.
 - **Resize and Pin NUMA** - Resizes the virtual CPU and NUMA topology of the virtual machine according to the Host, and pins them to the Host resources.
 - **Dedicated** - Exclusively pins virtual CPUs to host physical CPUs. Available for cluster compatibility level 4.7 or later. If the virtual machine has NUMA enabled, all nodes must be unpinned.
 - **Isolate Threads** - Exclusively pins virtual CPUs to host physical CPUs. Each virtual CPU gets a physical core. Available for cluster compatibility level 4.7 or later. If the virtual machine has NUMA enabled, all nodes must be unpinned.
10. Click **OK**.

**NOTE**

If you do not pin the virtual NUMA node to a host NUMA node, the system defaults to the NUMA node that contains the host device's memory-mapped I/O (MMIO), provided that there are one or more host devices and all of those devices are from a single NUMA node.


4.8. CONFIGURING SATELLITE ERRATA VIEWING FOR A VIRTUAL MACHINE

In the Administration Portal, you can configure a virtual machine to display the available errata. The virtual machine needs to be associated with a Red Hat Satellite server to show available errata.

Red Hat Virtualization 4.4 supports viewing errata with Red Hat Satellite 6.6.

Prerequisites

- The Satellite server must be added as an external provider.
- The Manager and any virtual machines on which you want to view errata must all be registered in the Satellite server by their respective FQDNs. This ensures that external content host IDs do not need to be maintained in Red Hat Virtualization.

**IMPORTANT**

Virtual machines added using an IP address cannot report errata.

- The host that the virtual machine runs on also needs to be configured to receive errata information from Satellite.
- The virtual machine must have the **ovirt-guest-agent** package installed. This package enables the virtual machine to report its host name to the Red Hat Virtualization Manager, which enables the Red Hat Satellite server to identify the virtual machine as a content host and report the applicable errata.
- The virtual machine must be registered to the Red Hat Satellite server as a content host.
- Use Red Hat Satellite remote execution to manage packages on hosts.

**NOTE**

The Katello agent is deprecated and will be removed in a future Satellite version. Migrate your processes to use the remote execution feature to update clients remotely.

Procedure

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Foreman/Satellite** tab.
4. Select the required Satellite server from the **Provider** drop-down list.
5. Click **OK**.

Additional resources

- [Setting up Satellite errata viewing for a host](#) in the *Administration Guide*
- [Installing the Guest Agents, Tools, and Drivers on Linux](#) in the *Virtual Machine Management Guide* for Red Hat Enterprise Linux virtual machines.
- [Installing the Guest Agents, Tools, and Drivers on Windows](#) in the *Virtual Machine Management Guide* for Windows virtual machines.

4.9. CONFIGURING HEADLESS VIRTUAL MACHINES

You can configure a headless virtual machine when it is not necessary to access the machine via a graphical console. This headless machine will run without graphical and video devices. This can be useful in situations where the host has limited resources, or to comply with virtual machine usage requirements such as real-time virtual machines.


Headless virtual machines can be administered via a Serial Console, SSH, or any other service for command line access. Headless mode is applied via the **Console** tab when creating or editing virtual machines and machine pools, and when editing templates. It is also available when creating or editing instance types.

If you are creating a new headless virtual machine, you can use the **Run Once** window to access the virtual machine via a graphical console for the first run only. See [Virtual Machine Run Once settings explained](#) for more details.

Prerequisites

- If you are editing an existing virtual machine, and the Red Hat Virtualization guest agent has not been installed, note the machine's IP prior to selecting **Headless Mode**.
- Before running a virtual machine in headless mode, the GRUB configuration for this machine must be set to console mode otherwise the guest operating system's boot process will hang. To set console mode, comment out the splashimage flag in the GRUB menu configuration file:

```
#splashimage=(hd0,0)/grub/splash.xpm.gz serial --unit=0 --speed=9600 --parity=no --stop=1
terminal --timeout=2 serial
```



NOTE

Restart the virtual machine if it is running when selecting the **Headless Mode** option.

Configuring a Headless Virtual Machine

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Console** tab.
4. Select **Headless Mode**. All other fields in the **Graphical Console** section are disabled.
5. Optionally, select **Enable VirtIO serial console** to enable communicating with the virtual machine via serial console. This is highly recommended.

6. Reboot the virtual machine if it is running. See [Rebooting a Virtual Machine](#).

4.10. CONFIGURING HIGH PERFORMANCE VIRTUAL MACHINES, TEMPLATES, AND POOLS

You can configure a virtual machine for high performance, so that it runs with performance metrics as close to bare metal as possible. When you choose high performance optimization, the virtual machine is configured with a set of automatic, and recommended manual, settings for maximum efficiency.

The high performance option is only accessible in the Administration Portal, by selecting **High Performance** from the **Optimized for** dropdown list in the **Edit** or **New** virtual machine, template, or pool window. This option is not available in the VM Portal.

The high performance option is supported by Red Hat Virtualization 4.2 and later. It is not available for earlier versions.

Virtual Machines

If you change the optimization mode of a running virtual machine to high performance, some configuration changes require restarting the virtual machine.

To change the optimization mode of a new or existing virtual machine to high performance, you may need to make manual changes to the cluster and to the pinned host configuration first.

A high performance virtual machine has certain limitations, because enhanced performance has a trade-off in decreased flexibility:

- If pinning is set for CPU threads, I/O threads, emulator threads, or NUMA nodes, according to the recommended settings, only a subset of cluster hosts can be assigned to the high performance virtual machine.
- Many devices are automatically disabled, which limits the virtual machine's usability.

Templates and Pools

High performance templates and pools are created and edited in the same way as virtual machines. If a high performance template or pool is used to create new virtual machines, those virtual machines inherits this property and its configurations. Certain settings, however, are not inherited and must be set manually:

- CPU pinning
- Virtual NUMA and NUMA pinning topology
- I/O and emulator threads pinning topology
- Pass-through Host CPU

4.10.1. Creating a High Performance Virtual Machine, Template, or Pool

To create a high performance virtual machine, template, or pool:

1. In the **New** or **Edit** window, select **High Performance** from the **Optimized for** drop-down menu.
Selecting this option automatically performs certain configuration changes to this virtual machine, which you can view by clicking different tabs. You can change them back to their

original settings or override them. (See [Automatic High Performance Configuration Settings](#) for details.) If you change a setting, its latest value is saved.

2. Click **OK**.

If you have not set any manual configurations, the **High Performance Virtual Machine/Pool Settings** screen describing the recommended manual configurations appears.

If you have set some of the manual configurations, the **High Performance Virtual Machine/Pool Settings** screen displays the settings you have not made.


If you have set all the recommended manual configurations, the **High Performance Virtual Machine/Pool Settings** screen does not appear.

3. If the **High Performance Virtual Machine/Pool Settings** screen appears, click **Cancel** to return to the **New** or **Edit** window to perform the manual configurations. See [Configuring the Recommended Manual Settings](#) for details.

Alternatively, click **OK** to ignore the recommendations. The result may be a drop in the level of performance.

4. Click **OK**.

You can view the optimization type in the **General** tab of the details view of the virtual machine, pool, or template.



NOTE

Certain configurations can override the high performance settings. For example, if you select an instance type for a virtual machine before selecting **High Performance** from the **Optimized for** drop-down menu and performing the manual configuration, the instance type configuration will not affect the high performance configuration. If, however, you select the instance type after the high performance configurations, you should verify the final configuration in the different tabs to ensure that the high performance configurations have not been overridden by the instance type.

The last-saved configuration usually takes priority.



NOTE

Support for instance types is now deprecated, and will be removed in a future release.

4.10.1.1. Automatic High Performance Configuration Settings

The following table summarizes the automatic settings. The **Enabled (Y/N)** column indicates configurations that are enabled or disabled. The **Applies to** column indicates the relevant resources:

- VM - Virtual machine
- T - Template
- P - Pool
- C - Cluster

Table 4.3. Automatic High Performance Configuration Settings

Setting	Enabled (Y/N)	Applies to
Headless Mode (Console tab)	Y	VM, T, P
USB Enabled (Console tab)	N	VM, T, P
Smartcard Enabled (Console tab)	N	VM, T, P
Soundcard Enabled (Console tab)	N	VM, T, P
Enable VirtIO serial console (Console tab)	Y	VM, T, P
Allow manual migration only (Host tab)	Y	VM, T, P
Pass-Through Host CPU (Host tab)	Y	VM, T, P
Highly Available ^[1] (High Availability tab)	N	VM, T, P
No-Watchdog (High Availability tab)	N	VM, T, P
Memory Balloon Device (Resource Allocation tab)	N	VM, T, P
I/O Threads Enabled ^[2] (Resource Allocation tab)	Y	VM, T, P
Paravirtualized Random Number Generator PCI (virtio-rng) device (Random Generator tab)	Y	VM, T, P
I/O and emulator threads pinning topology	Y	VM, T
CPU cache layer 3	Y	VM, T, P

1. **Highly Available** is not automatically enabled. If you select it manually, high availability should be enabled for pinned hosts only.
2. Number of I/O threads = 1.

4.10.1.2. I/O and Emulator Threads Pinning Topology (Automatic Settings)

The I/O and emulator threads pinning topology is a new configuration setting for Red Hat Virtualization 4.2. It requires that I/O threads, NUMA nodes, and NUMA pinning be enabled and set for the virtual machine. Otherwise, a warning will appear in the engine log.

Pinning topology:

- The first two CPUs of each NUMA node are pinned.
- If all vCPUs fit into one NUMA node of the host:
 - The first two vCPUs are automatically reserved/pinned

- The remaining vCPUs are available for manual vCPU pinning
- If the virtual machine spans more than one NUMA node:
 - The first two CPUs of the NUMA node with the most pins are reserved/pinned
 - The remaining pinned NUMA node(s) are for vCPU pinning only

Pools do not support I/O and emulator threads pinning.









WARNING

If a host CPU is pinned to both a vCPU and I/O and emulator threads, a warning will appear in the log and you will be asked to consider changing the CPU pinning topology to avoid this situation.

4.10.1.3. High Performance Icons

The following icons indicate the states of a high performance virtual machine in the **Compute → Virtual Machines** screen.

Table 4.4. High Performance Icons

Icon	Description
	High performance virtual machine
	High performance virtual machine with Next Run configuration
	Stateless, high performance virtual machine
	Stateless, high performance virtual machine with Next Run configuration
	Virtual machine in a high performance pool
	Virtual machine in a high performance pool with Next Run configuration

4.10.2. Configuring the Recommended Manual Settings

You can configure the recommended manual settings in either the **New** or the **Edit** windows.

If a recommended setting is not performed, the **High Performance Virtual Machine/Pool Settings** screen displays the recommended setting when you save the resource.

The recommended manual settings are:

- [Pinning CPUs](#)

- [Setting the NUMA Pinning Policy](#)
- [Configuring Huge Pages](#)
- [Disabling KSM](#)

4.10.2.1. Manual High Performance Configuration Settings

The following table summarizes the recommended manual settings. The **Enabled (Y/N)** column indicates configurations that should be enabled or disabled. The **Applies to** column indicates the relevant resources:

- VM - Virtual machine
- T - Template
- P - Pool
- C - Cluster

Table 4.5. Manual High Performance Configuration Settings


Setting	Enabled (Y/N)	Applies to
NUMA Node Count (Host tab)	Y	VM
Tune Mode (NUMA Pinning screen)	Y	VM
NUMA Pinning (Host tab)	Y	VM
CPU Pinning topology (Resource Allocation tab)	Y	VM, P
hugepages (Custom Properties tab)	Y	VM, T, P
KSM (Optimization tab)	N	C

4.10.2.2. Pinning CPUs

To pin vCPUs to a specific host's physical CPU:

1. In the **Host** tab, select the **Specific Host(s)** radio button.
2. In the **Resource Allocation** tab, enter the **CPU Pinning Topology**, verifying that the configuration fits the pinned host's configuration. See [Virtual Machine Resource Allocation settings explained](#) for information about the syntax of this field.
This field is populated automatically and the CPU topology is updated when automatic NUMA pinning is activated.
3. Verify that the virtual machine configuration is compatible with the host configuration:
 - A virtual machine's number of sockets must not be greater than the host's number of sockets.

- A virtual machine's number of cores per virtual socket must not be greater than the host's number of cores.
- CPU-intensive workloads perform best when the host and virtual machine expect the same cache usage. To achieve the best performance, a virtual machine's number of threads per core must not be greater than that of the host.



IMPORTANT

CPU pinning has the following requirements:

- If the host is NUMA-enabled, the host's NUMA settings (memory and CPUs) must be considered because the virtual machine has to fit the host's NUMA configuration.
- The [I/O and emulator threads pinning topology](#) must be considered.
- CPU pinning can only be set for virtual machines and pools, but not for templates. Therefore, you must set CPU pinning manually whenever you create a high performance virtual machine or pool, even if they are based on a high performance template.

4.10.2.3. Setting the NUMA Pinning Policy

To set the NUMA Pinning Policy, you need a NUMA-enabled pinned host with at least two NUMA nodes.

To set the NUMA pinning policy manually:

1. Click **NUMA Pinning**.
2. In the **NUMA Topology** window, click and drag virtual NUMA nodes from the box on the right to the host's physical NUMA nodes on the left as required.
3. Select **Strict**, **Preferred**, or **Interleave** from the **Tune Mode** drop-down list in each NUMA node. If the selected mode is **Preferred**, the **NUMA Node Count** must be set to **1**.
4. Click **OK**.

To set the NUMA pinning policy automatically:

1. In the **Resource Allocation** tab, under **CPU Allocation**, select **Resize and Pin NUMA** from the **CPU Pinning Policy** drop-down list.
2. Click **OK**.



IMPORTANT

The number of declared virtual NUMA nodes and the NUMA pinning policy must take into account:

- The host's NUMA settings (memory and CPUs)
- The NUMA node in which the host devices are declared
- The CPU pinning topology
- The [IO and emulator threads pinning topology](#)
- Huge page sizes
- NUMA pinning can only be set for virtual machines, not for pools or templates. You must set NUMA pinning manually when you create a high performance virtual machine based on a template.

4.10.2.4. Configuring Huge Pages

Huge pages are pre-allocated when a virtual machine starts to run (dynamic allocation is disabled by default).

To configure huge pages:

1. In the **Custom Properties** tab, select **hugepages** from the custom properties list, which displays **Please select a key...** by default.
2. Enter the huge page size in KB.
You should set the huge page size to the largest size supported by the pinned host. The recommended size for x86_64 is 1 GiB.

The huge page size has the following requirements:

- The virtual machine's huge page size must be the same size as the pinned host's huge page size.
- The virtual machine's memory size must fit into the selected size of the pinned host's free huge pages.
- The NUMA node size must be a multiple of the huge page's selected size.

To enable dynamic allocation of huge pages:

1. Disable the HugePages filter in the scheduler.
2. In the **[performance]** section in `/etc/vdsm/vdsm.conf` set the following:
`use_dynamic_hugepages = true`

Comparison between dynamic and static hugepages

The following table outlines advantages and disadvantages of dynamic and static hugepages.

Table 4.6. Dynamic vs static hugepages

Setting	Advantages	Disadvantages	Recommendations
dynamic hugepages	<ul style="list-style-type: none"> • Less configuration required • Less wasted memory (i.e. hugepages free on a host waiting for possible incoming migrations) 	Failure to allocate due to fragmentation	Use 2MB hugepages
static hugepages	Predictable results	<ul style="list-style-type: none"> • Requires a kernel command line in the Edit Host configuration in the Administration Portal. See Custom kernel command line • Requires a host reboot. 	



IMPORTANT

The following limitations apply:

- Memory hotplug/unplug is disabled
- The host's memory resource is limited

4.10.2.5. Disabling KSM

To disable Kernel Same-page Merging (KSM) for the cluster:

1. Click **Compute** → **Clusters** and select the cluster.
2. Click **Edit**.
3. In the **Optimization** tab, clear the **Enable KSM** check box.

4.11. CONFIGURING THE TIME ZONE

Red Hat Virtualization stores time zone configurations for virtual machines in **/etc/ovirt-engine/conf/00-timezone.properties**. This file contains default time zone values such as **Etc/GMT=Greenwich Standard Time**. It features mappings that are valid for Windows and non-Windows time zones.



IMPORTANT

Do not edit the actual **00-timezone.properties** file. Changes will be overwritten if you upgrade or restore the Manager.

Do not change values that come directly from the operating system or the Manager.

Procedure

1. Create an override file in `/etc/ovirt-engine/conf/`. The file name must begin with a value greater than **00**, so that the file appears after `/etc/ovirt-engine/conf/00-timezone.properties`, and ends with the extension, **.properties**.
For example, **10-timezone.properties** overrides the default file, **00-timezone.properties**. The last file in the file list has precedence over earlier files.
2. Add new time zones to that file. Be sure each key is a valid General time zone from the [time zone database](#) and the value is a valid Windows time zone:

General

Time zones used for non-Windows operating system types, must follow the standard time zone format for example, **Etc/GMT** or **Asia/Jerusalem**.

Windows

Time zones specifically supported on [Windows](#) for example, **GMT Standard Time** or **Israel Standard Time**.

3. Restart the **rhvm** service:

```
# systemctl restart ovirt-engine
```

CHAPTER 5. EDITING VIRTUAL MACHINES

5.1. EDITING VIRTUAL MACHINE PROPERTIES

Changes to storage, operating system, or networking parameters can adversely affect the virtual machine. Ensure that you have the correct details before attempting to make any changes. Virtual machines can be edited while running, and some changes (listed in the procedure below) will be applied immediately. To apply all other changes, the virtual machine must be shut down and restarted.




NOTE

External virtual machines (marked with the prefix **external**) cannot be edited through the Red Hat Virtualization Manager.

Editing Virtual Machines

1. Click **Compute** → **Virtual Machines**.
2. Select the virtual machine to be edited.
3. Click **Edit**.
4. Change settings as required.
Changes to the following settings are applied immediately:
 - **Name**
 - **Description**
 - **Comment**
 - **Optimized for** (Desktop/Server/High Performance)
 - **Delete Protection**
 - **Network Interfaces**
 - **Memory Size** (Edit this field to hot plug virtual memory. See [Hot Plugging Virtual Memory](#).)
 - **Virtual Sockets** (Edit this field to hot plug CPUs. See [CPU hot plug](#).)
 - **Highly Available**
 - **Priority for Run/Migration queue**
 - **Disable strict user checking**
 - **Icon**
5. Click **OK**.
6. If the **Next Start Configuration** pop-up window appears, click **OK**.

Some changes are applied immediately. All other changes are applied when you shut down and restart your virtual machine. Until then, the pending changes icon () appears as a reminder to restart the virtual machine.

5.2. NETWORK INTERFACES

5.2.1. Adding a New Network Interface

You can add multiple network interfaces to virtual machines. Doing so allows you to put your virtual machine on multiple logical networks.



NOTE

You can create an overlay network for your virtual machines, isolated from the hosts, by defining a logical network that is not attached to the physical interfaces of the host. For example, you can create a DMZ environment, in which the virtual machines communicate among themselves over the bridge created in the host.

The overlay network uses OVN, which must be installed as an external network provider. See the [Administration Guide](#) for more information

Procedure

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Network Interfaces** tab.
4. Click **New**.
5. Enter the **Name** of the network interface.
6. Select the **Profile** and the **Type** of network interface from the drop-down lists. The **Profile** and **Type** drop-down lists are populated in accordance with the profiles and network types available to the cluster and the network interface cards available to the virtual machine.
7. Select the **Custom MAC address** check box and enter a MAC address for the network interface card as required.
8. Click **OK**.

The new network interface is listed in the **Network Interfaces** tab in the details view of the virtual machine. The **Link State** is set to **Up** by default when the network interface card is defined on the virtual machine and connected to the network.

For more details on the fields in the **New Network Interface** window, see [Virtual Machine Network Interface dialogue entries](#).

5.2.2. Editing a Network Interface

In order to change any network settings, you must edit the network interface. This procedure can be performed on virtual machines that are running, but some actions can be performed only on virtual machines that are not running.

Editing Network Interfaces

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Network Interfaces** tab and select the network interface to edit.
4. Click **Edit**.
5. Change settings as required. You can specify the **Name**, **Profile**, **Type**, and **Custom MAC address**. See [Adding a Network Interface](#).
6. Click **OK**.

5.2.3. Hot Plugging a Network Interface

You can hot plug network interfaces. Hot plugging means enabling and disabling devices while a virtual machine is running.



NOTE

The guest operating system must support hot plugging network interfaces.

Hot Plugging Network Interfaces

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.
3. Click the **Network Interfaces** tab and select the network interface to hot plug.
4. Click **Edit**.
5. Set the **Card Status** to **Plugged** to enable the network interface, or set it to **Unplugged** to disable the network interface.
6. Click **OK**.

5.2.4. Removing a Network Interface

Removing Network Interfaces

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Network Interfaces** tab and select the network interface to remove.
4. Click **Remove**.
5. Click **OK**.

5.2.5. Configuring a virtual machine to ignore NICs

You can configure the **ovirt-guest-agent** on a virtual machine to ignore certain NICs. This prevents IP addresses associated with network interfaces created by certain software from appearing in reports. You must specify the name and number of the network interface you want to ignore (for example, **eth0**, **docker0**).

Procedure

1. In the **/etc/ovirt-guest-agent.conf** configuration file on the virtual machine, insert the following line, with the NICs to be ignored separated by spaces:

```
ignored_nics = first_NIC_to_ignore second_NIC_to_ignore
```

2. Start the agent:

```
# systemctl start ovirt-guest-agent
```



NOTE

Some virtual machine operating systems automatically start the guest agent during installation.

If your virtual machine's operating system automatically starts the guest agent or if you need to configure the denylist on many virtual machines, use the configured virtual machine as a template for creating additional virtual machines. See [Creating a template from an existing virtual machine](#) for details.

5.3. VIRTUAL DISKS

5.3.1. Adding a New Virtual Disk

You can add multiple virtual disks to a virtual machine.

Image is the default type of disk. You can also add a **Direct LUN** disk. **Image** disk creation is managed entirely by the Manager. **Direct LUN** disks require externally prepared targets that already exist. Existing disks are either floating disks or shareable disks attached to virtual machines.

Adding Disks to Virtual Machines

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Disks** tab.
4. Click **New**.
5. Use the appropriate radio buttons to switch between **Image** and **Direct LUN**.
6. Enter a **Size(GB)**, **Alias**, and **Description** for the new disk.
7. Use the drop-down lists and check boxes to configure the disk. See [Add Virtual Disk dialogue entries](#) for more details on the fields for all disk types.
8. Click **OK**.

The new disk appears in the details view after a short time.

5.3.2. Attaching an Existing Disk to a Virtual Machine

Floating disks are disks that are not associated with any virtual machine.

Floating disks can minimize the amount of time required to set up virtual machines. Designating a floating disk as storage for a virtual machine makes it unnecessary to wait for disk preallocation at the time of a virtual machine's creation.

Floating disks can be attached to a single virtual machine, or to multiple virtual machines if the disk is shareable. Each virtual machine that uses the shared disk can use a different disk interface type.

Once a floating disk is attached to a virtual machine, the virtual machine can access it.

Procedure

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Disks** tab.
4. Click **Attach**.
5. Select one or more virtual disks from the list of available disks and select the required interface from the **Interface** drop-down.
6. Click **OK**.



NOTE

No Quota resources are consumed by attaching virtual disks to, or detaching virtual disks from, virtual machines.

5.3.3. Extending the Available Size of a Virtual Disk

You can extend the available size of a virtual disk while the virtual disk is attached to a virtual machine. Resizing a virtual disk does not resize the underlying partitions or file systems on that virtual disk. Use the **fdisk** utility to resize the partitions and file systems as required. See [How to Resize a Partition using fdisk](#) for more information.

Extending the Available Size of Virtual Disks

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Disks** tab and select the disk to edit.
4. Click **Edit**.
5. Enter a value in the **Extend size by(GB)** field.
6. Click **OK**.

The target disk's status becomes **locked** for a short time, during which the drive is resized. When the resizing of the drive is complete, the status of the drive becomes **OK**.

5.3.4. Hot Plugging a Virtual Disk


You can hot plug virtual disks. Hot plugging means enabling or disabling devices while a virtual machine is running.



NOTE


The guest operating system must support hot plugging virtual disks.

Hot Plugging Virtual Disks

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Disks** tab and select the virtual disk to hot plug.
4. Click **More Actions** (), then click **Activate** to enable the disk, or **Deactivate** to disable the disk.
5. Click **OK**.

5.3.5. Removing a Virtual Disk from a Virtual Machine

Removing Virtual Disks From Virtual Machines

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine name to go to the details view.
3. Click the **Disks** tab and select the virtual disk to remove.
4. Click **More Actions** (), then click **Deactivate**.
5. Click **OK**.
6. Click **Remove**.
7. Optionally, select the **Remove Permanently** check box to completely remove the virtual disk from the environment. If you do not select this option - for example, because the disk is a shared disk - the virtual disk will remain in **Storage** → **Disks**.
8. Click **OK**.

If the disk was created as block storage, for example iSCSI, and the **Wipe After Delete** check box was selected when creating the disk, you can view the log file on the host to confirm that the data has been wiped after permanently removing the disk. See [Settings to Wipe Virtual Disks After Deletion](#) in the *Administration Guide*.


If the disk was created as block storage, for example iSCSI, and the **Discard After Delete** check box was selected on the storage domain before the disk was removed, a **blkdiscard** command is called on the logical volume when it is removed and the underlying storage is notified that the blocks are free. See

[Setting Discard After Delete for a Storage Domain](#) in the *Administration Guide*. A **blkdiscard** is also called on the logical volume when a virtual disk is removed if the virtual disk is attached to at least one virtual machine with the **Enable Discard** check box selected.

5.3.6. Importing a Disk Image from an Imported Storage Domain

You can import floating virtual disks from an imported storage domain.

This procedure requires access to the Administration Portal.



NOTE

Only QEMU-compatible disks can be imported into the Manager.

Importing a Disk Image

1. Click **Storage** → **Domains**.
2. Click an imported storage domain to go to the details view.
3. Click **Disk Import**.
4. Select one or more disk images and click **Import**. This opens the **Import Disk(s)** window.
5. Select the appropriate **Disk Profile** for each disk.
6. Click **OK** to import the selected disks.

5.3.7. Importing an Unregistered Disk Image from an Imported Storage Domain

You can import floating virtual disks from a storage domain. Floating disks created outside of a Red Hat Virtualization environment are not registered with the Manager. Scan the storage domain to identify unregistered floating disks to be imported.


This procedure requires access to the Administration Portal.



NOTE

Only QEMU-compatible disks can be imported into the Manager.

Importing a Disk Image

1. Click **Storage** → **Domains**.
2. Click **More Actions** (), then click **Scan Disks** so that the Manager can identify unregistered disks.
3. Select an unregistered disk name and click **Disk Import**.
4. Select one or more disk images and click **Import**. This opens the **Import Disk(s)** window.
5. Select the appropriate **Disk Profile** for each disk.
6. Click **OK** to import the selected disks.

5.4. VIRTUAL MEMORY

5.4.1. Hot Plugging Virtual Memory

You can hot plug virtual memory. Hot plugging means enabling or disabling devices while a virtual machine is running. Each time memory is hot plugged, it appears as a new memory device in the **Vm Devices** tab in the details view of the virtual machine, up to a maximum of 16 available slots. When the virtual machine is restarted, these devices are cleared from the **Vm Devices** tab without reducing the virtual machine's memory, allowing you to hot plug more memory devices. If the hot plug fails (for example, if there are no more available slots), the memory increase will be applied when the virtual machine is restarted.



IMPORTANT


This feature is currently not supported for the self-hosted engine Manager virtual machine.



NOTE

If you might need to later hot unplug the memory that you are now hot plugging, see [Hot Unplugging Virtual Memory](#).

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Edit**.
3. Click the **System** tab.
4. Increase the **Memory Size** by entering the total amount required. Memory can be added in multiples of 256 MB. By default, the maximum memory allowed for the virtual machine is set to 4x the memory size specified. Though the value is changed in the user interface, the maximum value is not hot plugged, and you will see the pending changes icon (). To avoid that, you can change the maximum memory back to the original value.
5. Click **OK**.
This action opens the **Pending Virtual Machine changes** window, as some values such as **maxMemorySizeMb** and **minAllocatedMem** will not change until the virtual machine is restarted. However, the hot plug action is triggered by the change to the **Memory Size** value, which can be applied immediately.
6. Click **OK**.

The virtual machine's **Defined Memory** is updated in the **General** tab in the details view. You can see the newly added memory device in the **Vm Devices** tab in the details view.

5.4.2. Hot Unplugging Virtual Memory

You can hot unplug virtual memory. Hot unplugging disables devices while a virtual machine is running.

Prerequisites

- Only memory added with hot plugging can be hot unplugged.

- The virtual machine's operating system must support memory hot unplugging.
- The virtual machine must not have a memory balloon device enabled. This feature is disabled by default.
- All blocks of the hot-plugged memory must be set to **online_movable** in the virtual machine's device management rules. In virtual machines running up-to-date versions of Red Hat Enterprise Linux or CoreOS, this rule is set by default. For information on device management rules, consult the documentation for the virtual machine's operating system.
- To ensure that hot plugged memory can be hot unplugged later, add the **movable_node** option to the kernel command line of the virtual machine as follows and reboot the virtual machine:

```
# grubby --update-kernel=ALL --args="movable_node"
```

For more information, see [Setting kernel command-line parameters](#) in the RHEL 8 document *Managing, monitoring and updating the kernel*.

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click the **Vm Devices** tab.
3. In the **Hot Unplug** column, click **Hot Unplug** beside the memory device to be removed.
4. Click **OK** in the **Memory Hot Unplug** window.
The **Physical Memory Guaranteed** value for the virtual machine is decremented automatically if necessary.

5.5. HOT PLUGGING VCPUS

You can hot plug vCPUs. Hot plugging means enabling or disabling devices while a virtual machine is running.



IMPORTANT

Hot unplugging a vCPU is only supported if the vCPU was previously hot plugged. A virtual machine's vCPUs cannot be hot unplugged to less vCPUs than it was originally created with.

The following prerequisites apply:

- The virtual machine's **Operating System** must be explicitly set in the **New Virtual Machine** or **Edit Virtual Machine** window.
- The virtual machine's operating system must support CPU hot plug. See the table below for support details.
- Windows virtual machines must have the guest agents installed. See [Installing the Guest Agents and Drivers on Windows](#).

Hot Plugging vCPUs

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Edit**.
3. Click the **System** tab.
4. Change the value of **Virtual Sockets** as required.
5. Click **OK**.

Table 5.1. Operating System Support Matrix for vCPU Hot Plug

Operating System	Version	Architecture	Hot Plug Supported	Hot Unplug Supported
Red Hat Enterprise Linux Atomic Host 7		x86	Yes	Yes
Red Hat Enterprise Linux 6.3+		x86	Yes	Yes
Red Hat Enterprise Linux 7.0+		x86	Yes	Yes
Red Hat Enterprise Linux 7.3+		PPC64	Yes	Yes
Red Hat Enterprise Linux 8.0+		x86	Yes	Yes
Microsoft Windows Server 2012 R2	All	x64	Yes	No
Microsoft Windows Server 2016	Standard, Datacenter	x64	Yes	No
Microsoft Windows Server 2019	Standard, Datacenter	x64	Yes	No
Microsoft Windows 8.x	All	x86	Yes	No
Microsoft Windows 8.x	All	x64	Yes	No
Microsoft Windows 10	All	x86	Yes	No

Operating System	Version	Architecture	Hot Plug Supported	Hot Unplug Supported
Microsoft Windows 10	All	x64	Yes	No

5.6. PINNING A VIRTUAL MACHINE TO MULTIPLE HOSTS

Virtual machines can be pinned to multiple hosts. Multi-host pinning allows a virtual machine to run on a specific subset of hosts within a cluster, instead of one specific host or all hosts in the cluster. The virtual machine cannot run on any other hosts in the cluster even if all of the specified hosts are unavailable. Multi-host pinning can be used to limit virtual machines to hosts with, for example, the same physical hardware configuration.

If a host fails, a highly available virtual machine is automatically restarted on one of the other hosts to which the virtual machine is pinned.

Pinning Virtual Machines to Multiple Hosts

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Host** tab.
4. Select the **Specific Host(s)** radio button under **Start Running On** and select two or more hosts from the list.
5. Click the **High Availability** tab.
6. Select the **Highly Available** check box.
7. Select **Low**, **Medium**, or **High** from the **Priority** drop-down list. When migration is triggered, a queue is created in which the high priority virtual machines are migrated first. If a cluster is running low on resources, only the high priority virtual machines are migrated.
8. Click **OK**.

5.7. VIEWING VIRTUAL MACHINES PINNED TO A HOST

You can view virtual machines pinned to a host even while the virtual machines are offline. Use the **Pinned to Host** list to see which virtual machines will be affected and which virtual machines will require a manual restart after the host becomes active again.


Viewing Virtual Machines Pinned to a Host

1. Click **Compute** → **Hosts**.
2. Click a host name to go to the details view.
3. Click the **Virtual Machines** tab.
4. Click **Pinned to Host**.

5.8. CHANGING THE CD FOR A VIRTUAL MACHINE

You can change the CD accessible to a virtual machine while that virtual machine is running, using ISO images that have been uploaded to the data domain of the virtual machine's cluster. See [Uploading Images to a Data Storage Domain](#) in the *Administration Guide* for details.

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **More Actions** (), then click **Change CD**.
3. Select an option from the drop-down list:
 - Select an ISO file from the list to eject the CD currently accessible to the virtual machine and mount that ISO file as a CD. .Procedure from the list to eject the CD currently accessible to the virtual machine.
4. Click **OK**.

5.9. SMART CARD AUTHENTICATION

Smart cards are an external hardware security feature, most commonly seen in credit cards, but also used by many businesses as authentication tokens. Smart cards can be used to protect Red Hat Virtualization virtual machines.

Enabling Smart Cards

1. Ensure that the smart card hardware is plugged into the client machine and is installed according to manufacturer's directions.
2. Click **Compute** → **Virtual Machines** and select a virtual machine.
3. Click **Edit**.
4. Click the **Console** tab and select the **Smartcard enabled** check box.
5. Click **OK**.
6. Connect to the running virtual machine by clicking the **Console** button. Smart card authentication is now passed from the client hardware to the virtual machine.



IMPORTANT

If the Smart card hardware is not correctly installed, enabling the Smart card feature will result in the virtual machine failing to load properly.

Disabling Smart Cards

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Console** tab, and clear the **Smartcard enabled** check box.

4. Click **OK**.

Configuring Client Systems for Smart Card Sharing

- Smart cards may require certain libraries in order to access their certificates. These libraries must be visible to the NSS library, which **spice-gtk** uses to provide the smart card to the guest. NSS expects the libraries to provide the PKCS #11 interface.
- Make sure that the module architecture matches the **spice-gtk** / **remote-viewer** architecture. For instance, if you have only the 32b PKCS #11 library available, you must install the 32b build of virt-viewer in order for smart cards to work.

Configuring RHEL Clients for Smart Card support

Red Hat Enterprise Linux provides support for Smart cards. Install the **Smart card support** group. If the Smart Card Support group is installed on a Red Hat Enterprise Linux system, smart cards are redirected to the guest when Smart Cards are enabled.

1. To install the **Smart card support** group, run the following command:

```
# dnf groupinstall "Smart card support"
```

Configuring RHEL Clients with Other Smart Card Middleware

Red Hat Enterprise Linux provides a system-wide registry of pkcs11 modules in the **p11-kit**, and these are accessible to all applications.

1. To register the third party PKCS#11 library in the p11-kit database, run the following command as root:

```
# echo "module: /path/to/library.so" > /etc/pkcs11/modules/my.module
```

2. To verify the Smart card is visible for p11-kit through this library run the following command:

```
$ p11-kit list-modules
```

Configuring Windows Clients

Red Hat does not provide PKCS #11 support to Windows clients. Libraries that provide PKCS #11 support must be obtained from third parties.

1. When such libraries are obtained, register them by running the following command as a user with elevated privileges:

```
modutil -dbdir %PROGRAMDATA%\pki\nssdb -add "module name" -libfile  
C:\Path\to\module_.dll
```

CHAPTER 6. ADMINISTRATIVE TASKS

6.1. SHUTTING DOWN A VIRTUAL MACHINE

You can turn off a virtual machine using **Shutdown** or **Power Off**. **Shutdown** gracefully shuts down a virtual machine. **Power Off** executes a hard shutdown. A graceful shutdown is usually preferable to a hard shutdown.



NOTE

If an exclamation point appears next to the virtual machine, a snapshot deletion process has failed, and you may not be able to restart the machine after shutting it down. Try to delete the snapshot again and ensure that the explanation mark disappears before shutting down the virtual machine. See [Deleting a snapshot](#) for more information.

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Shutdown** or right-click the virtual machine and select **Shutdown** from the pop-up menu.
3. Optionally in the Administration Portal, enter a **Reason** for shutting down the virtual machine in the **Shut down Virtual Machine(s)** confirmation window. This allows you to provide an explanation for the shutdown, which will appear in the logs and when the virtual machine is powered on again.



NOTE

The virtual machine shutdown **Reason** field will only appear if it has been enabled in the cluster settings. For more information, see [Explanation of Settings and Controls in the New Cluster and Edit Cluster Windows](#) in the *Administration Guide*.

4. Click **OK** in the **Shut down Virtual Machine(s)** confirmation window.

If the virtual machine gracefully shuts down, the **Status** of the virtual machine changes to **Down**. If the virtual machine does not gracefully shut down, click the down arrow next to **Shutdown** and then click **Power Off** to execute a hard shutdown, or right-click the virtual machine and select **Power Off** from the pop-up menu.

6.2. SUSPENDING A VIRTUAL MACHINE

Suspending a virtual machine is equal to placing that virtual machine into **Hibernate** mode.

Suspending a Virtual Machine

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Suspend** or right-click the virtual machine and select **Suspend** from the pop-up menu.


The **Status** of the virtual machine changes to **Suspended**.

6.3. REBOOTING OR RESETTING A VIRTUAL MACHINE

You can restart your virtual machines in two different ways; either using reboot or reset.

Several situations can occur where you need to reboot the virtual machine, such as after an update or configuration change. When you reboot, the virtual machine's console remains open while the guest operating system is restarted.

If a guest operating system can not be loaded or has become unresponsive, you need to reset the virtual machine. When you reset, the virtual machine's console remains open while the guest operating system is restarted.



NOTE

The reset operation can only be performed from the Administration Portal.

Rebooting a Virtual Machine

To reboot a virtual machine:

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Reboot** or right-click the virtual machine and select **Reboot** from the pop-up menu.
3. Click **OK** in the **Reboot Virtual Machine(s)** confirmation window.

Resetting a Virtual Machine

To reset a virtual machine:

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click the down arrow next to **Reboot**, then click **Reset**, or right-click the virtual machine and select **Reset** from the pop-up menu.
3. Click **OK** in the **Reset Virtual Machine(s)** confirmation window.

During reboot and reset operations, the **Status** of the virtual machine changes to **Reboot In Progress** before returning to **Up**.

6.4. REMOVING A VIRTUAL MACHINE



IMPORTANT

The **Remove** button is disabled while virtual machines are running; you must shut down a virtual machine before you can remove it.

Removing Virtual Machines


1. Click **Compute** → **Virtual Machines** and select the virtual machine to remove.
2. Click **Remove**.
3. Optionally, select the **Remove Disk(s)** check box to remove the virtual disks attached to the virtual machine together with the virtual machine. If the **Remove Disk(s)** check box is cleared, then the virtual disks remain in the environment as floating disks.

4. Click **OK**.

6.5. CLONING A VIRTUAL MACHINE

You can clone virtual machines without having to create a template or a snapshot first.

Procedure

1. Click **Compute** → **Virtual Machines** and select the virtual machine to clone.
2. Click **More Actions** (), then click **Clone VM**.
3. Enter a **Clone Name** for the new virtual machine.
4. Click **OK**.

6.6. UPDATING VIRTUAL MACHINE GUEST AGENTS AND DRIVERS

The Red Hat Virtualization guest agents, tools, and drivers provide additional functionality for virtual machines, such as gracefully shutting down or rebooting virtual machines from the VM Portal and Administration Portal. The tools and agents also provide information for virtual machines, including:

- Resource usage
- IP addresses
- Installed applications

The guest tools are distributed as an ISO file that you can attach to virtual machines. This ISO file is packaged as an RPM file that you can install and update from the Manager machine.

6.6.1. Updating the Guest Agents and Drivers on Red Hat Enterprise Linux

Update the guest agents and drivers on your Red Hat Enterprise Linux virtual machines to use the latest version.

Updating the Guest Agents and Drivers on Red Hat Enterprise Linux

1. Log in to the Red Hat Enterprise Linux virtual machine.
2. Update the **ovirt-guest-agent-common** package:

```
# yum update ovirt-guest-agent-common
```

3. Restart the service:

- For Red Hat Enterprise Linux 6

```
# service ovirt-guest-agent restart
```

- For Red Hat Enterprise Linux 7

```
# systemctl restart ovirt-guest-agent.service
```

6.6.2. Updating Windows drivers with Windows Update

When you need to update the drivers for a Windows virtual machine, the simplest method is to use Windows Update.

Procedure

1. Log in to the virtual machine.
2. Ensure that Windows Update is enabled so you can get updates.
3. Check Windows Update for updates from **Red Hat, Inc**
4. Manually install any updates that have not been automatically installed.

Additional resources

- [Updating Windows guest agents and drivers using the command prompt](#)
- See the Microsoft documentation for details on using Windows Update.

6.6.3. Updating Windows guest agents and drivers using the command prompt

When you do not have access to Windows Update to update Windows drivers, or when you need to update the oVirt guest agents, you can do so from the **virtio-win** package by using the virtual machine's command prompt. During this procedure, you must remove and reinstall the drivers, which can lead to network disruption. This procedure restores your settings after reinstalling the drivers.

Procedure

1. If you are updating the drivers, on the Windows virtual machine, use the **netsh** utility to save TCP settings before uninstalling the netkvm driver:

```
C:\WINDOWS\system32>netsh dump > filename.txt
```

2. On the Manager machine, update the **virtio-win** package to the latest version:

```
# dnf upgrade -y virtio-win
```

The **virtio-win_ version.iso** file is located in **/usr/share/virtio-win/** on the Manager machine.

3. Upload the ISO file to a data domain. For more information, see [Uploading Images to a Data Storage Domain](#) in the *Administration Guide*.
4. In the Administration or VM Portal, if the virtual machine is running, use the **Change CD** drop-down list to attach the **virtio-win_ version.iso** file to each of your virtual machines. If the virtual machine is powered off, click the **Run Once** button and attach the ISO as a CD.
5. Log in to the virtual machine.
6. Select the CD Drive (**D:** for this example) containing the **virtio-win_ version.iso** file.
7. Reinstall the guest agents or drivers:
 - To reinstall only the guest agents, use **qemu-ga-x86_64.msi**:

–

```
C:\WINDOWS\system32>msiexec.exe /i D:\guest-agent\qemu-ga-x86_64.msi /passive /norestart
```

- To reinstall the drivers, use **virtio-win-gt-x64.msi**:

```
C:\WINDOWS\system32>msiexec.exe /i D:\virtio-win-gt-x64.msi /passive /norestart
```

8. If you are updating the drivers, restore the settings you saved using **netsh**:

```
C:\WINDOWS\system32>netsh -f filename.txt
```

6.7. VIEWING RED HAT SATELLITE ERRATA FOR A VIRTUAL MACHINE

Errata for each virtual machine can be viewed after the Red Hat Virtualization virtual machine has been configured to receive errata information from the Red Hat Satellite server.

For more information on configuring a virtual machine to display available errata see [Configuring Satellite Errata](#)

Viewing Red Hat Satellite Errata

1. Click **Compute** → **Virtual Machines**.
2. Click the virtual machine's name to go to the details view.
3. Click **Errata**.

6.8. VIRTUAL MACHINES AND PERMISSIONS

6.8.1. Managing System Permissions for a Virtual Machine

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A **UserVmManager** is a system administration role for virtual machines in a data center. This role can be applied to specific virtual machines, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual resources.

The user virtual machine administrator role permits the following actions:

- Create, edit, and remove virtual machines.
- Run, suspend, shutdown, and stop virtual machines.



NOTE

You can only assign roles and permissions to existing users.

Many end users are concerned solely with the virtual machine resources of the virtualized environment. As a result, Red Hat Virtualization provides several user roles which enable the user to manage virtual machines specifically, but not other resources in the data center.

6.8.2. Virtual Machine Administrator Roles Explained

The table below describes the administrator roles and privileges applicable to virtual machine administration.

Table 6.1. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
DataCenterAdmin	Data Center Administrator	Possesses administrative permissions for all objects underneath a specific data center except for storage.
ClusterAdmin	Cluster Administrator	Possesses administrative permissions for all objects underneath a specific cluster.
NetworkAdmin	Network Administrator	Possesses administrative permissions for all operations on a specific logical network. Can configure and manage networks attached to virtual machines. To configure port mirroring on a virtual machine network, apply the NetworkAdmin role on the network and the UserVmManager role on the virtual machine.

6.8.3. Virtual Machine User Roles Explained

The table below describes the user roles and privileges applicable to virtual machine users. These roles allow access to the VM Portal for managing and accessing virtual machines, but they do not confer any permissions for the Administration Portal.

Table 6.2. Red Hat Virtualization System User Roles

Role	Privileges	Notes
UserRole	Can access and use virtual machines and pools.	Can log in to the VM Portal and use virtual machines and pools.

Role	Privileges	Notes
PowerUserRole	Can create and manage virtual machines and templates.	Apply this role to a user for the whole environment with the Configure window, or for specific data centers or clusters. For example, if a PowerUserRole is applied on a data center level, the PowerUser can create virtual machines and templates in the data center. Having a PowerUserRole is equivalent to having the VmCreator , DiskCreator , and TemplateCreator roles.
UserVmManager	System administrator of a virtual machine.	Can manage virtual machines and create and use snapshots. A user who creates a virtual machine in the VM Portal is automatically assigned the UserVmManager role on the machine.
UserTemplateBasedVm	Limited privileges to only use Templates.	Level of privilege to create a virtual machine by means of a template.
VmCreator	Can create virtual machines in the VM Portal.	This role is not applied to a specific virtual machine; apply this role to a user for the whole environment with the Configure window. When applying this role to a cluster, you must also apply the DiskCreator role on an entire data center, or on specific storage domains.
VnicProfileUser	Logical network and network interface user for virtual machines.	If the Allow all users to use this Network option was selected when a logical network is created, VnicProfileUser permissions are assigned to all users for the logical network. Users can then attach or detach virtual machine network interfaces to or from the logical network.

6.8.4. Assigning Virtual Machines to Users

If you are creating virtual machines for users other than yourself, you have to assign roles to the users before they can use the virtual machines. Note that permissions can only be assigned to existing users. See [Users and Roles](#) in the *Administration Guide* for details on creating user accounts.

The VM Portal supports three default roles: User, PowerUser and UserVmManager. However, customized roles can be configured via the Administration Portal. The default roles are described below.

- A **User** can connect to and use virtual machines. This role is suitable for desktop end users performing day-to-day tasks.
- A **PowerUser** can create virtual machines and view virtual resources. This role is suitable if you are an administrator or manager who needs to provide virtual resources for your employees.
- A **UserVmManager** can edit and remove virtual machines, assign user permissions, use snapshots and use templates. It is suitable if you need to make configuration changes to your virtual environment.

When you create a virtual machine, you automatically inherit **UserVmManager** privileges. This enables you to make changes to the virtual machine and assign permissions to the users you manage, or users who are in your Identity Management (IdM) or RHDS group. See the [Administration Guide](#) for more information.

Procedure

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.
3. Click the **Permissions** tab.
4. Click **Add**.
5. Enter a name, or user name, or part thereof in the **Search** text box, and click **Go**. A list of possible matches display in the results list.
6. Select the check box of the user to be assigned the permissions.
7. Select **UserRole** from the **Role to Assign** drop-down list.
8. Click **OK**.

The user's name and role display in the list of users permitted to access this virtual machine.



NOTE

If a user is assigned permissions to only one virtual machine, single sign-on (SSO) can be configured for the virtual machine. With single sign-on enabled, when a user logs in to the VM Portal, and then connects to a virtual machine through, for example, a SPICE console, users are automatically logged in to the virtual machine and do not need to type in the user name and password again. Single sign-on can be enabled or disabled on a per virtual machine basis. See [Configuring Single Sign-On for Virtual Machines](#) for more information on how to enable and disable single sign-on for virtual machines.

6.8.5. Removing Access to Virtual Machines from Users

Removing Access to Virtual Machines from Users

1. Click **Compute** → **Virtual Machines**.
2. Click the virtual machine's name to go to the details view.

3. Click **Permissions**.
4. Click **Remove**. A warning message displays, asking you to confirm removal of the selected permissions.
5. To proceed, click **OK**. To abort, click **Cancel**.

6.9. SNAPSHOTS

6.9.1. Creating a Snapshot of a Virtual Machine

A snapshot is a view of a virtual machine's operating system and applications on any or all available disks at a given point in time. Take a snapshot of a virtual machine before you make a change to it that may have unintended consequences. You can use a snapshot to restore a virtual machine to a previous state.


Creating a Snapshot of a Virtual Machine

In the VM Portal:

1. Open a virtual machine.
2. In the **Snapshots** panel, click **+Create Snapshot**.
A snapshot is added to the panel, including all attached disks.

In the Administration Portal:

1. Click **Compute** → **Virtual Machines**.
2. Click a virtual machine's name to go to the details view.
3. Click the **Snapshots** tab and click **Create**.
4. Enter a description for the snapshot.
5. Select **Disks to include** using the check boxes.



NOTE

If no disks are selected, a partial snapshot of the virtual machine, without a disk, is created. You can preview this snapshot to view the configuration of the virtual machine. Note that committing a partial snapshot will result in a virtual machine without a disk.

6. Select **Save Memory** to include a running virtual machine's memory in the snapshot.
7. Click **OK**.

The virtual machine's operating system and applications on the selected disk(s) are stored in a snapshot that can be previewed or restored. The snapshot is created with a status of **Locked**, which changes to **Ok**. When you click the snapshot, its details are shown on the **General**, **Disks**, **Network Interfaces**, and **Installed Applications** drop-down views in the **Snapshots** tab.

6.9.2. Using a Snapshot to Restore a Virtual Machine

A snapshot can be used to restore a virtual machine to its previous state.

Using Snapshots to Restore Virtual Machines

In the VM Portal:

1. Shutdown the virtual machine.
2. In the **Snapshots** panel, click the **Restore Snapshot** icon for the snapshot you want to restore. The snapshot is loaded.

In the Administration Portal:

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the name of the virtual machine to go to the details view.
3. Shut down the virtual machine.
4. Click the **Snapshots** tab to list the available snapshots.
5. Select a snapshot to restore in the upper pane. The snapshot details display in the lower pane.
6. Click the **Preview drop-down menu** button and select **Custom**.
7. Use the check boxes to select the **VM Configuration**, **Memory**, and disk(s) you want to restore, then click **OK**. This allows you to create and restore from a customized snapshot using the configuration and disk(s) from multiple snapshots.
The status of the snapshot changes to **Preview Mode**. The status of the virtual machine briefly changes to **Image Locked** before returning to **Down**.
8. Start the virtual machine; it runs using the disk image of the snapshot.
9. Click **Commit** to permanently restore the virtual machine to the condition of the snapshot. Any subsequent snapshots are erased.
Alternatively, click the **Undo** button to deactivate the snapshot and return the virtual machine to its previous state.

6.9.3. Creating a Virtual Machine from a Snapshot

You can use a snapshot to create another virtual machine.

Creating a Virtual Machine from a Snapshot

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.
3. Click the **Snapshots** tab to list the available snapshots.
4. Select a snapshot in the list displayed and click **Clone**.
5. Enter the **Name** of the virtual machine.
6. Click **OK**.

After a short time, the cloned virtual machine appears in the **Virtual Machines** tab in the navigation pane with a status of **Image Locked**. The virtual machine remains in this state until Red Hat Virtualization completes the creation of the virtual machine. A virtual machine with a preallocated 20 GB hard drive takes about fifteen minutes to create. Sparsely-allocated virtual disks take less time to create than do preallocated virtual disks.

When the virtual machine is ready to use, its status changes from **Image Locked** to **Down** in **Compute** → **Virtual Machines**.

6.9.4. Deleting a Snapshot

You can delete a virtual machine snapshot and permanently remove it from your Red Hat Virtualization environment.

Deleting a Snapshot

In the VM Portal:

1. Open a virtual machine.
2. In the **Snapshots** panel, click the **Delete Snapshot** icon of the snapshot you want to delete.

In the Administration Portal:

1. Click **Compute** → **Virtual Machines**.
2. Click the virtual machine's name to go to the details view.
3. Click the **Snapshots** tab to list the snapshots for that virtual machine.
4. Select the snapshot to delete.
5. Click **Delete**.
6. Click **OK**.



NOTE

If the deletion fails, fix the underlying problem (for example, a failed host, an inaccessible storage device, or a temporary network issue) and try again.

6.10. HOST DEVICES

6.10.1. Adding a Host Device to a Virtual Machine

To improve performance, you can attach a host device to a virtual machine.

Host devices are physical devices connected to a particular host machine, such as:

- SCSI tape drives, disks, and changers
- PCI NICs, GPUs, and HBAs
- USB mice, cameras, and disks

To add a host device to a virtual machine, you use the virtual machine's **Host Devices** properties. First, you select one of the cluster hosts and a device type. Then, you choose and attach one or more of the host devices on that host.



IMPORTANT

When you change the **Pinned Host** setting, it removes the current host devices. When you save these changes, in the virtual machine's **Host** settings, it sets **Start Running On** to **Specific Host(s)** and specifies the host you selected earlier using the **Pinned Host** setting.

When you finish attaching one or more host devices, you run the virtual machine to apply the changes. The virtual machine starts on the host that has the attached host devices.

If the virtual machine cannot start on the specified host or access the host device, it cancels the start operation and produces an error message with information about the cause.

Prerequisites

- The state of the host is **Up**.
- The host is configured for direct device assignment.

Procedure

1. In the Administration Portal, click **Compute → Virtual Machines**.
2. Shut down the virtual machine.
3. Click the name of the virtual machine to go to the details view.
4. Click the **Host Devices** tab.
5. Click **Add device**. This opens the **Add Host Devices** pane.
6. Use **Pinned Host** to select the host where the virtual machine runs.
7. Use **Capability** to list **pci**, **scsi**, **nvdimm**, or **usb_device** devices.



NOTE

The **nvdimm** option is a technical preview feature. For more information, see [nvdimm host devices](#).

8. Use **Available Host Devices** to select devices.
9. Click the down arrow to move devices to **Host Devices to be attached**
10. Click **OK** to attach these devices to the virtual machine and close the window.
11. Optional: If you attach a SCSI host device, configure the optimal driver.
 - a. Click the **Edit** button. This opens the **Edit Virtual Machine** pane.
 - b. Click the **Custom Properties** tab.

- c. Click the **Please select a key** and select **scsi_hostdev** from the bottom of the drop-down list.
 - d. In most cases, select **scsi-hd**. Otherwise, for tape or CD changer devices, select the **scsi_generic** option. For more details, see [Virtual Machine Custom Properties Settings Explained](#).
 - e. Click the **OK** button.
12. Run the virtual machine.
 13. While the virtual machine starts running, watch for **Operation Canceled** error messages.

Troubleshooting

If you cannot add a host device to a virtual machine, or a virtual machine cannot start running with the attached host devices, it generates **Operation Canceled** error messages. For example:

Operation Canceled

Error while executing action:

<vm name>:

- * Cannot run VM. There is no host that satisfies current scheduling constraints. See below for details:
- * The host <first_hostname> did not satisfy internal filter HostDevice because it does not support host device passthrough.
- * The host <second_hostname> did not satisfy internal filter HostDevice because the host does not provide requested host devices.

You can fix the error by removing the host device from the virtual machine or correcting the issues the error message describes. For example:

- Respond to a **The host <hostname> did not satisfy internal filter HostDevice because it does not support host device passthrough** message by configuring the host for device passthrough and restarting the virtual machine.
- Respond to the **The host <hostname> did not satisfy internal filter HostDevice because the host does not provide requested host devices** message by adding the host device to the host.
- Respond to a **Cannot add Host devices because the VM is in Up status** message by shutting down the virtual machine before adding a host device.
- Verify that the state of the host is **Up**.

Additional resources

- [Host Devices](#) in the *Virtual Machine Management Guide*.
- [Pinning a Virtual Machine to Multiple Hosts](#)
- [Configuring a Host for PCI Passthrough](#)
- [Additional Hardware Considerations for Using Device Assignment](#) in *Hardware Considerations for Implementing SR-IOV*.
- [nvdimms](#) host devices

6.10.2. Removing Host Devices from a Virtual Machine

If you are removing all host devices directly attached to the virtual machine in order to add devices from a different host, you can instead add the devices from the desired host, which will automatically remove all of the devices already attached to the virtual machine.

Procedure

1. Click **Compute** → **Virtual Machines**.
2. Select a virtual machine to go to the details view.
3. Click the **Host Devices** tab to list the host devices attached to the virtual machine.
4. Select the host device to detach from the virtual machine, or hold **Ctrl** to select multiple devices, and click **Remove device**. This opens the **Remove Host Device(s)** window.
5. Click **OK** to confirm and detach these devices from the virtual machine.

6.10.3. Pinning a Virtual Machine to Another Host

You can use the **Host Devices** tab in the details view of a virtual machine to pin it to a specific host.

If the virtual machine has any host devices attached to it, pinning it to another host automatically removes the host devices from the virtual machine.

Pinning a Virtual Machine to a Host

1. Click a virtual machine name and click the **Host Devices** tab.
2. Click **Pin to another host**. This opens the **Pin VM to Host** window.
3. Use the **Host** drop-down menu to select a host.
4. Click **OK** to pin the virtual machine to the selected host.

6.10.4. NVDIMM host devices



NOTE

NVDIMM devices are Technology Preview features only. Technology Preview features are not supported with Red Hat production service-level agreements (SLAs) and might not be functionally complete, and Red Hat does not recommend using them for production. These features provide early access to upcoming product features, enabling customers to test functionality and provide feedback during the development process. For more information, see [Red Hat Technology Preview Features Support Scope](#).

You can add emulated NVDIMM devices to virtual machines. Elsewhere, this type of memory is also known as *virtual NVDIMM* or *vNVDIMM*.


The emulated NVDIMM you can attach to a virtual machine is backed by real NVDIMM on the host machine where the virtual machine runs. Therefore, when you attach NVDIMM to a virtual machine, you also pin the virtual machine to a specific host.

You can reconfigure the mode, partitioning, and other properties of the emulated NVDIMM device in the virtual machine without affecting the settings of the physical NVDIMM on the host device.

To add emulated NVDIMM to a virtual machine, see [Adding Host Devices to a Virtual Machine](#)

Limitations

- Memory snapshots are disabled when an NVDIMM device is present in a virtual machine. There is no way to make a snapshot of NVDIMM content, and a memory snapshot cannot work correctly without having the corresponding NVDIMM data.
- In RHV, each NVDIMM device passed to a virtual machine has an automatically-assigned label area with a fixed size of 128 KB. IBM POWER hardware, and 128 KB is the minimum label size allowed by QEMU.
- By default, the virtual machine uses the whole NVDIMM device. You cannot configure the size of the NVDIMM from the virtual machine. To configure its size, partition the NVDIMM device on the host and add the partition to the virtual machine.
- The size of the NVDIMM device on the virtual machine may be slightly lower than on the host to comply with libvirt and QEMU alignment and size adjustments. Precise sizing is also needed to make memory hotplug work.
- libvirt and QEMU adjust their size and label placement. If those internal arrangements change, it can cause data loss.
- NVDIMM hotplug is not supported by the platform.
- Virtual machines with NVDIMM devices cannot migrate because they are pinned to a host.
- SELinux currently prevents access to NVDIMM devices in **devdax** mode. As a result, data persistence cannot be guaranteed if the host fails. See [BZ1855336](#).



IMPORTANT

Avoid using NVDIMM on IBM POWER hardware. This combination is currently not expected to be stable until further work is completed.

Additional resources

- [QEMU NVDIMM documentation](#)
- [Adding Host Devices to a Virtual Machine](#)

6.11. AFFINITY GROUPS

6.11.1. Affinity Groups

You can create **Affinity Groups** to help determine where selected virtual machines run in relation to each other and to specified hosts. This capability helps manage workload scenarios such as licensing requirements, high-availability workloads, and disaster recovery.

The VM Affinity Rule

When you create an **Affinity Group**, you select the virtual machines that belong to the group. To define

where these virtual machines can run in relation to each other, you enable a **VM Affinity Rule**: A **Positive** affinity rule tries to run the virtual machines together on a single host; a **Negative** affinity rule tries to run the virtual machines on separate hosts. If the rule cannot be fulfilled, the outcome depends on whether the weight or filter module is enabled.

The Host Affinity Rule

Optionally, you can add hosts to the **Affinity Group**. To define where virtual machines in the group can run in relation to hosts in the group, you enable a **Host Affinity Rule**: A **Positive** affinity rule tries to run the virtual machines on hosts in the affinity group; a **Negative** affinity rule tries to run the virtual machines on hosts that are not in the affinity group. If the rule cannot be fulfilled, the outcome depends on whether the weight or filter module is enabled.

The Default Weight Module

By default, both rules apply the *weight module* in the cluster's scheduling policy. With the weight module, the scheduler attempts to fulfill a rule, but allows the virtual machines in the affinity group to run anyway if the rule cannot be fulfilled.

For example, with a positive **VM Affinity Rule** and the weight module enabled, the scheduler *tries* to run all of the affinity group's virtual machines on a single host. However, if a single host does not have sufficient resources for this, the scheduler runs the virtual machines on multiple hosts.

For this module to work, the **weight module** section of the scheduling policies must contain the **VmAffinityGroups** and **VmToHostsAffinityGroups** keywords.

The Enforcing Option and Filter Module

Both rules have an **Enforcing** option which applies the *filter module* in the cluster's scheduling policy. The filter module overrides the weight module. With the filter module enabled, the scheduler *requires* that a rule be fulfilled. If a rule cannot be fulfilled, the filter module prevents the virtual machines in the affinity group from running.

For example, with a **Positive Host Affinity Rule** and **Enforcing** enabled (the filter module enabled), the scheduler *requires* the virtual machines in the affinity group to run on hosts that are part of the affinity group. However, if those hosts are down, the scheduler does not run the virtual machines at all.

For this module to work, the **filter module** section of the scheduling policies must contain the **VmAffinityGroups** and **VmToHostsAffinityGroups** keywords.

Examples

To see how these rules and options can be used with one another, see [Affinity group examples](#).



WARNING

- For affinity labels to work, the **filter module** section of the scheduling policies must contain **Label**.
- If an affinity group and affinity label conflict with each other, the affected virtual machines do not run. To help prevent, troubleshoot, and resolve conflicts, see [Affinity group troubleshooting](#).



IMPORTANT

Each rule is affected by the weight and filter modules in the cluster's scheduling policy.

- For the **VM Affinity Rule** rule to work, the scheduling policy must have the **VmAffinityGroups** keyword in its **Weight module** and **Filter module** sections.
- For the **Host Affinity Rule** to work, the scheduling policy must have the **VmToHostsAffinityGroups** keyword in its **Weight module** and **Filter module** sections.

For more information, see [Scheduling Policies](#) in the *Administration Guide*.



NOTE

- Affinity groups apply to virtual machines in a cluster. Moving a virtual machine from one cluster to another removes it from the affinity groups in the original cluster.
- Virtual machines do not have to restart for the affinity group rules to take effect.

6.11.2. Creating an Affinity Group

You can create new affinity groups in the Administration Portal.

Creating Affinity Groups

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.
3. Click the **Affinity Groups** tab.
4. Click **New**.
5. Enter a **Name** and **Description** for the affinity group.
6. From the **VM Affinity Rule** drop-down, select **Positive** to apply positive affinity or **Negative** to apply negative affinity. Select **Disable** to disable the affinity rule.
7. Select the **Enforcing** check box to apply hard enforcement, or ensure this check box is cleared to apply soft enforcement.
8. Use the drop-down list to select the virtual machines to be added to the affinity group. Use the + and - buttons to add or remove additional virtual machines.
9. Click **OK**.

6.11.3. Editing an Affinity Group

Editing Affinity Groups

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.

3. Click the **Affinity Groups** tab.
4. Click **Edit**.
5. Change the **VM Affinity Rule** drop-down and **Enforcing** check box to the preferred values and use the + and - buttons to add or remove virtual machines to or from the affinity group.
6. Click **OK**.

6.11.4. Removing an Affinity Group

Removing Affinity Groups

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the virtual machine's name to go to the details view.
3. Click the **Affinity Groups** tab.
4. Click **Remove**.
5. Click **OK**.

The affinity policy that applied to the virtual machines that were members of that affinity group no longer applies.

6.11.5. Affinity Groups Examples

The following examples illustrate how to apply affinity rules for various scenarios, using the different features of the affinity group capability described in this chapter.

Example 6.1. High Availability

Dalia is the DevOps engineer for a startup. For high availability, a particular system's two virtual machines should run on separate hosts anywhere in the cluster.

Dalia creates an affinity group named "high availability" and does the following:

- Adds the two virtual machines, **VM01** and **VM02**, to the affinity group.
- Sets **VM Affinity** to **Negative** so the virtual machines try to run on separate hosts.
- Leaves **Enforcing** cleared (disabled) so that both virtual machines can continue running in case only one host is available during an outage.
- Leaves the **Hosts** list empty so the virtual machines run on any host in the cluster.

Example 6.2. Performance

Sohni is a software developer who uses two virtual machines to build and test his software many times each day. There is heavy network traffic between these two virtual machines. Running the machines on the same host reduces both network traffic and the effects of network latency on the build and test process. Using high-specification hosts (faster CPUs, SSDs, and more memory) further accelerates this process.

Sohni creates an affinity group called "build and test" and does the following:

- Adds **VM01** and **VM02**, the build and test virtual machines, to the affinity group.
- Adds the high-specification hosts, **host03**, **host04**, and **host05**, to the affinity group.
- Sets **VM affinity** to **Positive** so the virtual machines try to run on the same host, reducing network traffic and latency effects.
- Sets **Host affinity** to **Positive** so the virtual machines try to run on the high specification hosts, accelerating the process.
- Leaves **Enforcing** cleared (disabled) for both rules so the virtual machines can run if the high-specification hosts are not available.

Example 6.3. Licensing

Bandile, a software asset manager, helps his organization comply with the restrictive licensing requirements of a 3D imaging software vendor. These terms require the virtual machines for its licensing server, **VM-LS**, and imaging workstations, **VM-WS#**, to run on the same host. Additionally, the physical CPU-based licensing model requires that the workstations run on either of two GPU-equipped hosts, **host-gpu-primary** or **host-gpu-backup**.

To meet these requirements, Bandile creates an affinity group called "3D seismic imaging" and does the following:

- Adds the previously mentioned virtual machines and hosts to the affinity group.
- Sets **VM affinity** to **Positive** and selects **Enforcing** so the licensing server and workstations *must* run together on one of the hosts, not on multiple hosts.
- Sets **Host affinity** to **Positive** and selects **Enforcing** so the virtual machines *must* run on either of the GPU-equipped the hosts, not other hosts in the cluster.

6.11.6. Affinity Groups Troubleshooting

To help prevent problems with affinity groups

- Plan and document the scenarios and outcomes you expect when using affinity groups.
- Verify and test the outcomes under a range of conditions.
- Follow change management best practices.
- Only use the **Enforcing** option if it is required.

If you observe problems with virtual machines not running

- Verify that the cluster has a scheduling policy with weight module and filter module sections that contain **VmAffinityGroups** and **VmToHostsAffinityGroups**. For more information, see [Explanation of Settings in the New Scheduling Policy and Edit Scheduling Policy Window](#) in the *Administration Guide*.

- Check for conflicts between affinity labels and affinity groups.

For possible conflicts between affinity labels and affinity groups

- If an affinity label and affinity group conflict with each other, the intersecting set of virtual machines do not run.
- To determine whether a conflict is possible:
 - Inspect the filter module section of the cluster's scheduling policies. These must contain both a **Label** keyword and a **VmAffinityGroups** OR **VmToHostsAffinityGroups** keyword. Otherwise, a conflict is *not possible*. (The presence of **VmAffinityGroups** and **VmToHostsAffinityGroups** in the weight module section does not matter because **Label** in a filter module section would override them.)
 - Inspect the affinity groups. They must contain a rule that has **Enforcing** enabled. Otherwise, a conflict is *not possible*.
- If a conflict is possible, identify the set of virtual machines that might be involved:
 - Inspect the affinity labels and groups. Make a list of virtual machines that are members of both an affinity label *and* an affinity group with an **Enforcing** option enabled.
 - For each host and virtual machine in this intersecting set, analyze the conditions under which a potential conflict occurs.
- Determine whether the actual non-running virtual machines match the ones in the analysis.
- Restructure the affinity groups and affinity labels to help avoid unintended conflicts.
- Verify that any changes produce the expected results under a range of conditions.
- If you have overlapping affinity groups and affinity labels, it can be easier to view them in one place as affinity groups. Consider converting an affinity label into an equivalent affinity group, which has a **Host affinity rule** with **Positive** selected and **Enforcing** enabled.

6.12. AFFINITY LABELS

6.12.1. About Affinity Labels

You can create and modify **Affinity Labels** in the Administration Portal.

Affinity Labels are used together with **Affinity Groups** to set any kind of affinity between virtual machines and hosts (hard, soft, positive, negative). See the [Affinity Groups](#) section for more information about affinity hardness and polarity.



WARNING

Affinity labels are a subset of affinity groups and can conflict with them. If there is a conflict, the virtual machine will not start.

6.12.2. Creating an Affinity Label

You can create affinity labels from the details view of a virtual machine, host, or cluster. This procedure uses the cluster details view.

Creating an Affinity Label

1. Click **Compute** → **Clusters** and select the appropriate cluster.
2. Click the cluster's name to go to the details view.
3. Click the **Affinity Labels** tab.
4. Click **New**.
5. Enter a **Name** for the affinity label.
6. Use the drop-down lists to select the virtual machines and hosts to be associated with the label. Use the + button to add additional virtual machines and hosts.
7. Click **OK**.

6.12.3. Editing an Affinity Label

You can edit affinity labels from the details view of a virtual machine, host, or cluster. This procedure uses the cluster details view.

Editing an Affinity Label

1. Click **Compute** → **Clusters** and select the appropriate cluster.
2. Click the cluster's name to go to the details view.
3. Click the **Affinity Labels** tab.
4. Select the label you want to edit.
5. Click **Edit**.
6. Use the + and - buttons to add or remove virtual machines and hosts to or from the affinity label.
7. Click **OK**.

6.12.4. Deleting an Affinity Label

You can only remove an **Affinity Label** from the details view of a cluster after it is deleted from each entity.

Deleting an Affinity Label

1. Click **Compute** → **Clusters** and select the appropriate cluster.
2. Click the cluster's name to go to the details view.
3. Click the **Affinity Labels** tab.

4. Select the label you want to remove.
5. Click **Edit**.
6. Use the - buttons to remove all virtual machines and hosts from the label.
7. Click **OK**.
8. Click **Delete**.
9. Click **OK**.

6.13. EXPORTING AND IMPORTING VIRTUAL MACHINES AND TEMPLATES



NOTE


The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center. See the [Importing Existing Storage Domains](#) section in the *Red Hat Virtualization Administration Guide* for information on importing storage domains.

You can export virtual machines and templates from, and import them to, data centers in the same or different Red Hat Virtualization environment. You can export or import virtual machines by using an export domain, a data domain, or by using a Red Hat Virtualization host.

When you export or import a virtual machine or template, properties including basic details such as the name and description, resource allocation, and high availability settings of that virtual machine or template are preserved.

The permissions and user roles of virtual machines and templates are included in the OVF files, so that when a storage domain is detached from one data center and attached to another, the virtual machines and templates can be imported with their original permissions and user roles. In order for permissions to be registered successfully, the users and roles related to the permissions of the virtual machines or templates must exist in the data center before the registration process.

You can also use the V2V feature to import virtual machines from other virtualization providers, such as RHEL 5 Xen or VMware, or import Windows virtual machines. V2V converts virtual machines so that they can be hosted by Red Hat Virtualization. For more information on installing and using V2V, see [Converting Virtual Machines from Other Hypervisors to KVM with virt-v2v](#).




IMPORTANT

Virtual machines must be shut down before being imported.

6.13.1. Exporting a Virtual Machine to the Export Domain

Export a virtual machine to the export domain so that it can be imported into a different data center. Before you begin, the export domain must be attached to the data center that contains the virtual machine to be exported.

Exporting a Virtual Machine to the Export Domain

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **More Actions** (), then click **Export to Export Domain**.
3. Optionally, select the following check boxes in the **Export Virtual Machine** window:
 - **Force Override**: overrides existing images of the virtual machine on the export domain.
 - **Collapse Snapshots**: creates a single export volume per disk. This option removes snapshot restore points and includes the template in a template-based virtual machine, and removes any dependencies a virtual machine has on a template. For a virtual machine that is dependent on a template, either select this option, export the template with the virtual machine, or make sure the template exists in the destination data center.



NOTE

When you create a virtual machine from a template by clicking **Compute** → **Templates** and clicking **New VM**, you will see two storage allocation options in the **Storage Allocation** section in the **Resource Allocation** tab:

- If **Clone** is selected, the virtual machine is not dependent on the template. The template does not have to exist in the destination data center.
- If **Thin** is selected, the virtual machine is dependent on the template, so the template must exist in the destination data center or be exported with the virtual machine. Alternatively, select the **Collapse Snapshots** check box to collapse the template disk and virtual disk into a single disk.

To check which option was selected, click a virtual machine's name and click the **General** tab in the details view.

4. Click **OK**.

The export of the virtual machine begins. The virtual machine displays in **Compute** → **Virtual Machines** with an **Image Locked** status while it is exported. Depending on the size of your virtual machine hard disk images, and your storage hardware, this can take up to an hour. Click the **Events** tab to view progress. When complete, the virtual machine has been exported to the export domain and displays in the **VM Import** tab of the export domain's details view.

6.13.2. Exporting a Virtual Machine to a Data Domain

You can export a virtual machine to a data domain to store a clone of the virtual machine as a backup.

When you export a virtual machine that is dependent on a template, the target storage domain should include that template.



NOTE

When you create a virtual machine from a template, you can choose from either of two storage allocation options:

- **Clone:** The virtual machine is not dependent on the template. The template does not have to exist in the destination storage domain.
- **Thin:** The virtual machine is dependent on the template, so the template must exist in the destination storage domain.

To check which option is selected, click a virtual machine's name and click the **General** tab in the details view.

Prerequisites

- The data domain is attached to a data center.
- The virtual machine is powered off.

Procedure

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Export**.
3. Specify a name for the exported virtual machine.
4. Select a target storage domain from the **Storage domain** pop-up menu.
5. (Optional) Check **Collapse snapshots** to export the virtual machine without any snapshots.
6. Click **OK**.

The Manager clones the virtual machine, including all its disks, to the target domain.



NOTE

When you move a disk from one type of data domain another, the disk format changes accordingly. For example, if the disk is on an NFS data domain, and it is in sparse format, then if you move the disk to an iSCSI domain its format changes to preallocated. This is different from using an export domain, because an export domain is NFS.

The virtual machine appears with an **Image Locked** status while it is exported. Depending on the size of your virtual machine hard disk images, and your storage hardware, this can take up to an hour. Click the **Events** tab to view progress. When complete, the virtual machine has been exported to the data domain and appears in the list of virtual machines.

Additional resources

- [Creating a Virtual Machine Based on a Template](#) in the *Virtual Machine Management Guide*

6.13.3. Importing a Virtual Machine from the Export Domain


You have a virtual machine on an export domain. Before the virtual machine can be imported to a new data center, the export domain must be attached to the destination data center.

Importing a Virtual Machine into the Destination Data Center

1. Click **Storage** → **Domains** and select the export domain. The export domain must have a status of **Active**.
2. Click the export domain's name to go to the details view.
3. Click the **VM Import** tab to list the available virtual machines to import.
4. Select one or more virtual machines to import and click **Import**.
5. Select the **Target Cluster**.
6. Select the **Collapse Snapshots** check box to remove snapshot restore points and include templates in template-based virtual machines.
7. Click the virtual machine to be imported and click the **Disks** sub-tab. From this tab, you can use the **Allocation Policy** and **Storage Domain** drop-down lists to select whether the disk used by the virtual machine will be thinly provisioned or preallocated, and can also select the storage domain on which the disk will be stored. An icon is also displayed to indicate which of the disks to be imported acts as the boot disk for that virtual machine.
8. Click **OK** to import the virtual machines.
The **Import Virtual Machine Conflict** window opens if the virtual machine exists in the virtualized environment.

Choose one of the following radio buttons:

- **Don't import**
 - **Import as cloned** and enter a unique name for the virtual machine in the **New Name** field.
9. Optionally select the **Apply to all** check box to import all duplicated virtual machines with the same suffix, and then enter a suffix in the **Suffix to add to the cloned VMs** field.
 10. Click **OK**.



IMPORTANT

During a single import operation, you can only import virtual machines that share the same architecture. If any of the virtual machines to be imported have a different architecture to that of the other virtual machines to be imported, a warning will display and you will be prompted to change your selection so that only virtual machines with the same architecture will be imported.

6.13.4. Importing a Virtual Machine from a Data Domain


You can import a virtual machine into one or more clusters from a data storage domain.

Prerequisite

- If you are importing a virtual machine from an imported data storage domain, the imported storage domain must be attached to a data center and activated.

Procedure

1. Click **Storage** → **Domains**.
2. Click the imported storage domain's name. This opens the details view.
3. Click the **VM Import** tab.
4. Select one or more virtual machines to import.
5. Click **Import**.
6. For each virtual machine in the **Import Virtual Machine(s)** window, ensure the correct target cluster is selected in the **Cluster** list.
7. Map external virtual machine vNIC profiles to profiles that are present on the target cluster(s):
 - a. Click **vNic Profiles Mapping**.
 - b. Select the vNIC profile to use from the **Target vNic Profile** drop-down list.
 - c. If multiple target clusters are selected in the **Import Virtual Machine(s)** window, select each target cluster in the **Target Cluster** drop-down list and ensure the mappings are correct.
 - d. Click **OK**.
8. If a MAC address conflict is detected, an exclamation mark appears next to the name of the virtual machine. Mouse over the icon to view a tooltip displaying the type of error that occurred. Select the **Reassign Bad MACs** check box to reassign new MAC addresses to all problematic virtual machines. Alternatively, you can select the **Reassign** check box per virtual machine.



NOTE

If there are no available addresses to assign, the import operation will fail. However, in the case of MAC addresses that are outside the cluster's MAC address pool range, it is possible to import the virtual machine without reassigning a new MAC address.

9. Click **OK**.

The imported virtual machines no longer appear in the list under the **VM Import** tab.

6.13.5. Importing a Virtual Machine from a VMware Provider

Import virtual machines from a VMware vCenter provider to your Red Hat Virtualization environment. You can import from a VMware provider by entering its details in the **Import Virtual Machine(s)** window during each import operation, or you can add the VMware provider as an external provider, and select the preconfigured provider during import operations. To add an external provider, see [Adding a VMware Instance as a Virtual Machine Provider](#).

Red Hat Virtualization uses V2V to import VMware virtual machines. For OVA files, the only disk format Red Hat Virtualization supports is VMDK.

**NOTE**

The **virt-v2v** package is not available on the ppc64le architecture and these hosts cannot be used as proxy hosts.

**NOTE**

If the import fails, refer to the relevant log file in **/var/log/vdsm/import/** and to **/var/log/vdsm/vdsm.log** on the proxy host for details.

Prerequisites

- The **virt-v2v** package must be installed on at least one host, referred to in this procedure as the proxy host. The **virt-v2v** package is available by default on Red Hat Virtualization Hosts and is installed on Red Hat Enterprise Linux hosts as a dependency of VDSM when added to the Red Hat Virtualization environment.
- Red Hat Enterprise Linux hosts must be Red Hat Enterprise Linux 7.2 or later.
- At least one data and one ISO storage domain are connected to the data center.


**NOTE**

You can only migrate to shared storage, such as NFS, iSCSI, or FCP. Local storage is not supported.

Although the ISO storage domain has been deprecated, it is required for migration.

- The **virtio-win_ version.iso** image file for Windows virtual machines is uploaded to the ISO storage domain. This image includes the guest tools that are required for migrating Windows virtual machines.
- The virtual machine must be shut down before being imported. Starting the virtual machine through VMware during the import process can result in data corruption.
- An import operation can only include virtual machines that share the same architecture. If any virtual machine to be imported has a different architecture, a warning appears and you are prompted to change your selection to include only virtual machines with the same architecture.

Procedure

1. Click **Compute** → **Virtual Machines**.
2. Click **More Actions** () and select **Import**. This opens the **Import Virtual Machine(s)** window.
3. Select **VMware** from the **Source** list.
4. If you have configured a VMware provider as an external provider, select it from the **External Provider** list. Verify that the provider credentials are correct. If you did not specify a destination data center or proxy host when configuring the external provider, select those options now.
5. If you have not configured a VMware provider, or want to import from a new VMware provider, provide the following details:

- a. Select from the list the **Data Center** in which the virtual machine will be available.
 - b. Enter the IP address or fully qualified domain name of the VMware vCenter instance in the **vCenter** field.
 - c. Enter the IP address or fully qualified domain name of the host from which the virtual machines will be imported in the **ESXi** field.
 - d. Enter the name of the data center and the cluster in which the specified ESXi host resides in the **Data Center** field.
 - e. If you have exchanged the SSL certificate between the ESXi host and the Manager, leave **Verify server's SSL certificate** checked to verify the ESXi host's certificate. If not, clear the option.
 - f. Enter the **Username** and **Password** for the VMware vCenter instance. The user must have access to the VMware data center and ESXi host on which the virtual machines reside.
 - g. Select a host in the chosen data center with **virt-v2v** installed to serve as the **Proxy Host** during virtual machine import operations. This host must also be able to connect to the network of the VMware vCenter external provider.
6. Click **Load** to list the virtual machines on the VMware provider that can be imported.
 7. Select one or more virtual machines from the **Virtual Machines on Source** list, and use the arrows to move them to the **Virtual Machines to Import** list. Click **Next**.



NOTE

If a virtual machine's network device uses the driver type e1000 or rtl8139, the virtual machine will use the same driver type after it has been imported to Red Hat Virtualization.

If required, you can change the driver type to VirtIO manually after the import. To change the driver type after a virtual machine has been imported, see [Editing network interfaces](#). If the network device uses driver types other than e1000 or rtl8139, the driver type is changed to VirtIO automatically during the import. The **Attach VirtIO-drivers** option allows the VirtIO drivers to be injected to the imported virtual machine files so that when the driver is changed to VirtIO, the device will be properly detected by the operating system.

8. Select the **Cluster** in which the virtual machines will reside.
9. Select a **CPU Profile** for the virtual machines.
10. Select the **Collapse Snapshots** check box to remove snapshot restore points and include templates in template-based virtual machines.
11. Select the **Clone** check box to change the virtual machine name and MAC addresses, and clone all disks, removing all snapshots. If a virtual machine appears with a warning symbol beside its name or has a tick in the **VM in System** column, you must clone the virtual machine and change its name.
12. Click each virtual machine to be imported and click the **Disks** sub-tab. Use the **Allocation Policy** and **Storage Domain** lists to select whether the disk used by the virtual machine will be thinly provisioned or preallocated, and select the storage domain on which the disk will be

stored. An icon is also displayed to indicate which of the disks to be imported acts as the boot disk for that virtual machine.

13. If you selected the **Clone** check box, change the name of the virtual machine in the **General** sub-tab.
14. Click **OK** to import the virtual machines.

The CPU type of the virtual machine must be the same as the CPU type of the cluster into which it is being imported. To view the cluster's **CPU Type** in the Administration Portal:

1. Click **Compute** → **Clusters**.
2. Select a cluster.
3. Click **Edit**.
4. Click the **General** tab.


If the CPU type of the virtual machine is different, configure the imported virtual machine's CPU type:

1. Click **Compute** → **Virtual Machines**.
2. Select the virtual machine.
3. Click **Edit**.
4. Click the **System** tab.
5. Click the **Advanced Parameters** arrow.
6. Specify the **Custom CPU Type** and click **OK**.

6.13.6. Exporting a Virtual Machine to a Host

You can export a virtual machine to a specific path or mounted NFS shared storage on a host in the Red Hat Virtualization data center. The export will produce an Open Virtual Appliance (OVA) package.

Exporting a Virtual Machine to a Host

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **More Actions** (), then click **Export to OVA**.
3. Select the host from the **Host** drop-down list.
4. Enter the absolute path to the export directory in the **Directory** field, including the trailing slash. For example: **/images2/ova/**
5. Optionally change the default name of the file in the **Name** field.
6. Click **OK**.

The status of the export can be viewed in the **Events** tab.

6.13.7. Importing a Virtual Machine from a Host

Import an Open Virtual Appliance (OVA) file into your Red Hat Virtualization environment. You can import the file from any Red Hat Virtualization Host in the data center.




IMPORTANT

Currently, only Red Hat Virtualization and OVAs created by VMware can be imported. KVM and Xen are not supported.

The import process uses **virt-v2v**. Only virtual machines running operating systems compatible with **virt-v2v** can be successfully imported. See [Converting virtual machines from other hypervisors to KVM with virt-v2v in RHEL 7 and RHEL 8](#) for a current list of compatible operating systems.

Importing an OVA File

1. Copy the OVA file to a host in your cluster, in a file system location such as `/var/tmp`.




NOTE

The location can be a local directory or a remote NFS mount, as long as it is not in the `/root`` directory or subdirectories. Ensure that it has sufficient space.

2. Ensure that the OVA file has permissions allowing read/write access to the **qemu** user (UID 36) and the **kvm** group (GID 36):

```
# chown 36:36 path_to_OVA_file/file.OVA
```

3. Click **Compute** → **Virtual Machines**.
4. Click **More Actions** () and select **Import**. This opens the **Import Virtual Machine(s)** window.
 - a. Select **Virtual Appliance (OVA)** from the **Source** list.
 - b. Select a host from the **Host** list.
 - c. In the **Path** field, specify the absolute path of the OVA file.
 - d. Click **Load** to list the virtual machine to be imported.
 - e. Select the virtual machine from the **Virtual Machines on Source** list, and use the arrows to move it to the **Virtual Machines to Import** list.
5. Click **Next**.
 - a. Select the **Storage Domain** for the virtual machine.
 - b. Select the **Target Cluster** where the virtual machines will reside.
 - c. Select the **CPU Profile** for the virtual machines.
 - d. Select the **Allocation Policy** for the virtual machines.
 - e. Optionally, select the **Attach VirtIO-Drivers** check box and select the appropriate image on the list to add VirtIO drivers.

- f. Select the **Allocation Policy** for the virtual machines.
 - g. Select the virtual machine, and on the **General** tab select the **Operating System**.
 - h. On the **Network Interfaces** tab, select the **Network Name** and **Profile Name**.
 - i. Click the **Disks** tab to view the **Alias**, **Virtual Size**, and **Actual Size** of the virtual machine.
6. Click **OK** to import the virtual machines.

6.13.8. Importing a virtual machine from a RHEL 5 Xen host

Import virtual machines from Xen on Red Hat Enterprise Linux 5 to your Red Hat Virtualization environment. Red Hat Virtualization uses V2V to import QCOW2 or raw virtual machine disk formats.

The **virt-v2v** package must be installed on at least one host (referred to in this procedure as the proxy host). The **virt-v2v** package is available by default on Red Hat Virtualization Hosts (RHVH) and is installed on Red Hat Enterprise Linux hosts as a dependency of VDSM when added to the Red Hat Virtualization environment. Red Hat Enterprise Linux hosts must be Red Hat Enterprise Linux 7.2 or later.



WARNING

If you are importing a Windows virtual machine from a RHEL 5 Xen host and you are using VirtIO devices, install the VirtIO drivers before importing the virtual machine. If the drivers are not installed, the virtual machine may not boot after import.


The VirtIO drivers can be installed from the **virtio-win_ version.iso** or the **RHV-toolsSetup_ version.iso**. See [Installing the Guest Agents and Drivers on Windows](#) for details.

If you are not using VirtIO drivers, review the configuration of the virtual machine before first boot to ensure that VirtIO devices are not being used.



NOTE

The **virt-v2v** package is not available on the ppc64le architecture and these hosts cannot be used as proxy hosts.



IMPORTANT

An import operation can only include virtual machines that share the same architecture. If any virtual machine to be imported has a different architecture, a warning appears and you are prompted to change your selection to include only virtual machines with the same architecture.



NOTE

If the import fails, refer to the relevant log file in **/var/log/vdsm/import/** and to **/var/log/vdsm/vdsm.log** on the proxy host for details.

Procedure

To import a virtual machine from RHEL 5 Xen, follow these steps:

1. Shut down the virtual machine. Starting the virtual machine through Xen during the import process can result in data corruption.
2. Enable public key authentication between the proxy host and the RHEL 5 Xen host:

- a. Log in to the proxy host and generate SSH keys for the **vds**m user.


```
# sudo -u vds m ssh-keygen
```


- b. Copy the **vds**m user's public key to the RHEL 5 Xen host.

```
# sudo -u vds m ssh-copy-id root@xenhost.example.com
```

- c. Log in to the RHEL 5 Xen host to verify that the login works correctly.

```
# sudo -u vds m ssh root@xenhost.example.com
```

3. Log in to the Administration Portal.
4. Click **Compute** → **Virtual Machines**.
5. Click **More Actions** () and select **Import**. This opens the **Import Virtual Machine(s)** window.
6. Select the **Data Center** that contains the proxy host.
7. Select **XEN (via RHEL)** from the **Source** drop-down list.
8. Optionally, select a RHEL 5 Xen **External Provider** from the drop-down list. The URI will be pre-filled with the correct URI. See [Adding a RHEL 5 Xen Host as a Virtual Machine Provider](#) in the *Administration Guide* for more information.
9. Enter the **URI** of the RHEL 5 Xen host. The required format is pre-filled; you must replace **<hostname>** with the host name of the RHEL 5 Xen host.
10. Select the proxy host from the **Proxy Host** drop-down list.
11. Click **Load** to list the virtual machines on the RHEL 5 Xen host that can be imported.
12. Select one or more virtual machines from the **Virtual Machines on Source** list, and use the arrows to move them to the **Virtual Machines to Import** list.



NOTE

Due to current limitations, Xen virtual machines with block devices do not appear in the **Virtual Machines on Source** list. They must be imported manually. See [Importing Block Based Virtual Machine from Xen host](#) .

13. Click **Next**.
14. Select the **Cluster** in which the virtual machines will reside.

15. Select a **CPU Profile** for the virtual machines.
16. Use the **Allocation Policy** and **Storage Domain** lists to select whether the disk used by the virtual machine will be thinly provisioned or preallocated, and select the storage domain on which the disk will be stored.

**NOTE**

The target storage domain must be a file-based domain. Due to current limitations, specifying a block-based domain causes the V2V operation to fail.

17. If a virtual machine appears with a warning symbol beside its name, or has a tick in the **VM in System** column, select the **Clone** check box to clone the virtual machine.

**NOTE**

Cloning a virtual machine changes its name and MAC addresses and clones all of its disks, removing all snapshots.

18. Click **OK** to import the virtual machines.

The CPU type of the virtual machine must be the same as the CPU type of the cluster into which it is being imported. To view the cluster's **CPU Type** in the Administration Portal:

1. Click **Compute → Clusters**.
2. Select a cluster.
3. Click **Edit**.
4. Click the **General** tab.

If the CPU type of the virtual machine is different, configure the imported virtual machine's CPU type:

1. Click **Compute → Virtual Machines**.
2. Select the virtual machine.
3. Click **Edit**.
4. Click the **System** tab.
5. Click the **Advanced Parameters** arrow.
6. Specify the **Custom CPU Type** and click **OK**.

Importing a Block-Based Virtual Machine from a RHEL 5 Xen Host

1. Enable public key authentication between the proxy host and the RHEL 5 Xen host:
 - a. Log in to the proxy host and generate SSH keys for the **vdsm** user.

```
# sudo -u vdsm ssh-keygen
```

- b. Copy the **vdsm** user's public key to the RHEL 5 Xen host.


```
# sudo -u vdsm ssh-copy-id root@xenhost.example.com
```

- c. Log in to the RHEL 5 Xen host to verify that the login works correctly.

```
# sudo -u vdsm ssh root@xenhost.example.com
```

2. Attach an export domain. See [Attaching an Existing Export Domain to a Data Center](#) in the *Administration Guide* for details.
3. On the proxy host, copy the virtual machine from the RHEL 5 Xen host:

```
# virt-v2v-copy-to-local -ic xen+ssh://root@xenhost.example.com vmname
```

4. Convert the virtual machine to libvirt XML and move the file to your export domain:

```
# virt-v2v -i libvirtxml vmname.xml -o rhev -of raw -os storage.example.com:/exportdomain
```

5. In the Administration Portal, click **Storage** → **Domains**, click the export domain's name, and click the **VM Import** tab in the details view to verify that the virtual machine is in your export domain.
6. Import the virtual machine into the destination data domain. See [Importing the virtual machine from the export domain](#) for details.

6.13.9. Importing a Virtual Machine from a KVM Host

Import virtual machines from KVM to your Red Hat Virtualization environment. Red Hat Virtualization converts KVM virtual machines to the correct format before they are imported. You must enable public key authentication between the KVM host and at least one host in the destination data center (this host is referred to in the following procedure as the proxy host).



WARNING

The virtual machine must be shut down before being imported. Starting the virtual machine through KVM during the import process can result in data corruption.



IMPORTANT

An import operation can only include virtual machines that share the same architecture. If any virtual machine to be imported has a different architecture, a warning appears and you are prompted to change your selection to include only virtual machines with the same architecture.



NOTE

If the import fails, refer to the relevant log file in `/var/log/vdsm/import/` and to `/var/log/vdsm/vdsm.log` on the proxy host for details.

Importing a Virtual Machine from KVM

1. Enable public key authentication between the proxy host and the KVM host:

- a. Log in to the proxy host and generate SSH keys for the **vds**m user.


```
# sudo -u vds m ssh-keygen
```

- b. Copy the **vds**m user's public key to the KVM host. The proxy host's **known_hosts** file will also be updated to include the host key of the KVM host.

```
# sudo -u vds m ssh-copy-id root@kvmhost.example.com
```

- c. Log in to the KVM host to verify that the login works correctly.

```
# sudo -u vds m ssh root@kvmhost.example.com
```

2. Log in to the Administration Portal.
3. Click **Compute** → **Virtual Machines**.
4. Click **More Actions** () and select **Import**. This opens the **Import Virtual Machine(s)** window.
5. Select the **Data Center** that contains the proxy host.
6. Select **KVM (via Libvirt)** from the **Source** drop-down list.
7. Optionally, select a KVM provider **External Provider** from the drop-down list. The URI will be pre-filled with the correct URI. See [Adding a KVM Host as a Virtual Machine Provider](#) in the *Administration Guide* for more information.
8. Enter the **URI** of the KVM host in the following format:

```
qemu+ssh://root@kvmhost.example.com/system
```

9. Keep the **Requires Authentication** check box selected.
10. Enter **root** in the **Username** field.
11. Enter the **Password** of the KVM host's root user.
12. Select the **Proxy Host** from the drop-down list.
13. Click **Load** to list the virtual machines on the KVM host that can be imported.
14. Select one or more virtual machines from the **Virtual Machines on Source** list, and use the arrows to move them to the **Virtual Machines to Import** list.
15. Click **Next**.
16. Select the **Cluster** in which the virtual machines will reside.
17. Select a **CPU Profile** for the virtual machines.
18. Optionally, select the **Collapse Snapshots** check box to remove snapshot restore points and include templates in template-based virtual machines.

19. Optionally, select the **Clone** check box to change the virtual machine name and MAC addresses, and clone all disks, removing all snapshots. If a virtual machine appears with a warning symbol beside its name or has a tick in the **VM in System** column, you must clone the virtual machine and change its name.
20. Click each virtual machine to be imported and click the **Disks** sub-tab. Use the **Allocation Policy** and **Storage Domain** lists to select whether the disk used by the virtual machine will be thin provisioned or preallocated, and select the storage domain on which the disk will be stored. An icon is also displayed to indicate which of the disks to be imported acts as the boot disk for that virtual machine. See [Virtual Disk Storage Allocation Policies](#) in the *Technical Reference* for more information.
21. If you selected the **Clone** check box, change the name of the virtual machine in the **General** tab.
22. Click **OK** to import the virtual machines.

The CPU type of the virtual machine must be the same as the CPU type of the cluster into which it is being imported. To view the cluster's **CPU Type** in the Administration Portal:

1. Click **Compute → Clusters**.
2. Select a cluster.
3. Click **Edit**.
4. Click the **General** tab.

If the CPU type of the virtual machine is different, configure the imported virtual machine's CPU type:

1. Click **Compute → Virtual Machines**.
2. Select the virtual machine.
3. Click **Edit**.
4. Click the **System** tab.
5. Click the **Advanced Parameters** arrow.
6. Specify the **Custom CPU Type** and click **OK**.

6.13.10. Importing a Red Hat KVM Guest Image

You can import a Red Hat-provided KVM virtual machine image. This image is a virtual machine snapshot with a preconfigured instance of Red Hat Enterprise Linux installed.

You can configure this image with the **cloud-init** tool, and use it to provision new virtual machines. This eliminates the need to install and configure the operating system and provides virtual machines that are ready for use.

Procedure

1. Download the most recent KVM virtual machine image from the [Download Red Hat Enterprise Linux](#) list, in the Product Software tab.
2. Upload the virtual machine image using the Manager or the REST API. See [Uploading Images to a Data Storage Domain](#) in the *Administration Guide*.

3. Create a new virtual machine and attach the uploaded disk image to it. See [Creating a Linux virtual machine](#).
4. Optionally, use **cloud-init** to configure the virtual machine. See [Using Cloud-Init to Automate the Configuration of Virtual Machines](#) for details.
5. Optionally, create a template from the virtual machine. You can generate new virtual machines from this template. See [Templates](#) for information about creating templates and generating virtual machines from templates.

6.14. MIGRATING VIRTUAL MACHINES BETWEEN HOSTS

Live migration provides the ability to move a running virtual machine between physical hosts with no interruption to service. The virtual machine remains powered on and user applications continue to run while the virtual machine is relocated to a new physical host. In the background, the virtual machine's RAM is copied from the source host to the destination host. Storage and network connectivity are not altered.



NOTE

A virtual machine that is using a vGPU cannot be migrated to a different host.

6.14.1. Live Migration Prerequisites




NOTE

This is one in a series of topics that show how to set up and configure SR-IOV on Red Hat Virtualization. For more information, see [Setting Up and Configuring SR-IOV](#)

You can use live migration to seamlessly move virtual machines to support a number of common maintenance tasks. Your Red Hat Virtualization environment must be correctly configured to support live migration well in advance of using it.

At a minimum, the following prerequisites must be met to enable successful live migration of virtual machines:

- The source and destination hosts are members of the same cluster, ensuring CPU compatibility between them.



NOTE

Live migrating virtual machines between different clusters is generally not recommended.

- The source and destination hosts' status is **Up**.
- The source and destination hosts have access to the same virtual networks and VLANs.
- The source and destination hosts have access to the data storage domain on which the virtual machine resides.
- The destination host has sufficient CPU capacity to support the virtual machine's requirements.
- The destination host has sufficient unused RAM to support the virtual machine's requirements.

- The migrating virtual machine does not have the **cache!=none** custom property set.

Live migration is performed using the management network and involves transferring large amounts of data between hosts. Concurrent migrations have the potential to saturate the management network. For best performance, create separate logical networks for management, storage, display, and virtual machine data to minimize the risk of network saturation.

6.14.2. Configuring Virtual Machines with SR-IOV-Enabled vNICs to Reduce Network Outage during Migration

Virtual machines with vNICs that are directly connected to a virtual function (VF) of an SR-IOV-enabled host NIC can be further configured to reduce network outage during live migration:

1. Ensure that the destination host has an available VF.
2. Set the **Passthrough** and **Migratable** options in the passthrough vNIC's profile. See [Enabling Passthrough on a vNIC Profile](#) in the *Administration Guide*.
3. Enable hotplugging for the virtual machine's network interface.
4. Ensure that the virtual machine has a backup VirtIO vNIC, in addition to the passthrough vNIC, to maintain the virtual machine's network connection during migration.
5. Set the VirtIO vNIC's **No Network Filter** option before configuring the bond. See [Explanation of Settings in the VM Interface Profile Window](#) in the *Administration Guide*.
6. Add both vNICs as slaves under an **active-backup** bond on the virtual machine, with the passthrough vNIC as the primary interface.

The bond and vNIC profiles can be configured in one of the following ways:

- The bond is not configured with **fail_over_mac=active** and the VF vNIC is the primary **slave** (recommended).
Disable the VirtIO vNIC profile's MAC-spoofing filter to ensure that traffic passing through the VirtIO vNIC is not dropped because it uses the VF vNIC MAC address.
- The bond is configured with **fail_over_mac=active**.
This failover policy ensures that the MAC address of the bond is always the MAC address of the active slave. During failover, the virtual machine's MAC address changes, with a slight disruption in traffic.

6.14.3. Configuring Virtual Machines with SR-IOV-Enabled vNICs with minimal downtime

To configure virtual machines for migration with SR-IOV enabled vNICs and minimal downtime follow the procedure described below.



NOTE

The following steps are provided only as a Technology Preview. For more information see [Red Hat Technology Preview Features Support Scope](#).

1. Create a vNIC profile with SR-IOV enabled vNICs. See [Creating a vNIC profile](#) and [Setting up and configuring SR-IOV](#).

2. In the Administration Portal, go to **Network → vNIC profiles**, select the vNIC profile, click **Edit** and select a **Failover vNIC profile** from the drop down list.
3. Click **OK** to save the profile settings.
4. Hotplug a network interface with the failover vNIC profile you created into the virtual machine, or start a virtual machine with this network interface plugged in.

**NOTE**

The virtual machine has three network interfaces: a controller interface and two secondary interfaces. The controller interface must be active and connected in order for migration to succeed.

5. For automatic deployment of virtual machines with this configuration, use the following **udev** rule:

```
UBSYSTEM=="net",
ACTION=="add|change",
ENV{ID_NET_DRIVER}!="net_failover",
ENV{NM_UNMANAGED}="1",
RUN+="/bin/sh -c '/sbin/ip link set up $INTERFACE'"
```

This **udev** rule works only on systems that manage interfaces with **NetworkManager**. This rule ensures that only the controller interface is activated.

6.14.4. Optimizing Live Migration

Live virtual machine migration can be a resource-intensive operation. To optimize live migration, you can set the following two options globally for every virtual machine in an environment, for every virtual machine in a cluster, or for an individual virtual machine.

**NOTE**

The **Auto Converge migrations** and **Enable migration compression** options are available for cluster levels 4.2 or earlier.

For cluster levels 4.3 or later, auto converge is enabled by default for all built-in migration policies, and migration compression is enabled by default for only the **Suspend workload if needed** migration policy. You can change these parameters when adding a new migration policy, or by modifying the **MigrationPolicies** configuration value.

The **Auto Converge migrations** option allows you to set whether auto-convergence is used during live migration of virtual machines. Large virtual machines with high workloads can dirty memory more quickly than the transfer rate achieved during live migration, and prevent the migration from converging. Auto-convergence capabilities in QEMU allow you to force convergence of virtual machine migrations. QEMU automatically detects a lack of convergence and triggers a throttle-down of the vCPUs on the virtual machine.

The **Enable migration compression** option allows you to set whether migration compression is used during live migration of the virtual machine. This feature uses Xor Binary Zero Run-Length-Encoding to reduce virtual machine downtime and total live migration time for virtual machines running memory write-intensive workloads or for any application with a sparse memory update pattern.

Both options are disabled globally by default.

Procedure

1. Enable auto-convergence at the global level:

```
# engine-config -s DefaultAutoConvergence=True
```

- a. Enable migration compression at the global level:

```
# engine-config -s DefaultMigrationCompression=True
```

- b. Restart the **ovirt-engine** service to apply the changes:

```
# systemctl restart ovirt-engine.service
```

2. Configure the optimization settings for a cluster:

- a. Click **Compute** → **Clusters** and select a cluster.
- b. Click **Edit**.
- c. Click the **Migration Policy** tab.
- d. From the **Auto Converge migrations** list, select **Inherit from global setting**, **Auto Converge**, or **Don't Auto Converge**.
- e. From the **Enable migration compression** list, select **Inherit from global setting**, **Compress**, or **Don't Compress**.
- f. Click **OK**.

3. Configure the optimization settings at the virtual machine level:

- a. Click **Compute** → **Virtual Machines** and select a virtual machine.
- b. Click **Edit**.
- c. Click the **Host** tab.
- d. From the **Auto Converge migrations** list, select **Inherit from cluster setting**, **Auto Converge**, or **Don't Auto Converge**.
- e. From the **Enable migration compression** list, select **Inherit from cluster setting**, **Compress**, or **Don't Compress**.
- f. Click **OK**.

6.14.5. Guest Agent Hooks

Hooks are scripts that trigger activity within a virtual machine when key events occur:

- Before migration
- After migration
- Before hibernation

- After hibernation

The hooks configuration base directory is **/etc/ovirt-guest-agent/hooks.d** on Linux systems.

Each event has a corresponding subdirectory: **before_migration** and **after_migration**, **before_hibernation** and **after_hibernation**. All files or symbolic links in that directory will be executed.

The executing user on Linux systems is **ovirtagent**. If the script needs **root** permissions, the elevation must be executed by the creator of the hook script.

6.14.6. Automatic Virtual Machine Migration

Red Hat Virtualization Manager automatically initiates live migration of all virtual machines running on a host when the host is moved into maintenance mode. The destination host for each virtual machine is assessed as the virtual machine is migrated, in order to spread the load across the cluster.

From version 4.3, all virtual machines defined with manual or automatic migration modes are migrated when the host is moved into maintenance mode. However, for high performance and/or pinned virtual machines, a **Maintenance Host** window is displayed, asking you to confirm the action because the performance on the target host may be less than the performance on the current host.

The Manager automatically initiates live migration of virtual machines in order to maintain load-balancing or power-saving levels in line with scheduling policy. Specify the scheduling policy that best suits the needs of your environment. You can also disable automatic, or even manual, live migration of specific virtual machines where required.

If your virtual machines are configured for high performance, and/or if they have been pinned (by setting Passthrough Host CPU, CPU Pinning, or NUMA Pinning), the migration mode is set to **Allow manual migration only**. However, this can be changed to **Allow Manual and Automatic** mode if required. Special care should be taken when changing the default migration setting so that it does not result in a virtual machine migrating to a host that does not support high performance or pinning.

6.14.7. Preventing Automatic Migration of a Virtual Machine

Red Hat Virtualization Manager allows you to disable automatic migration of virtual machines. You can also disable manual migration of virtual machines by setting the virtual machine to run only on a specific host.


The ability to disable automatic migration and require a virtual machine to run on a particular host is useful when using application high availability products, such as Red Hat High Availability or Cluster Suite.

Preventing Automatic Migration of Virtual Machines

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Host** tab.
4. In the **Start Running On** section, select **Any Host in Cluster** or **Specific Host(s)**, which enables you to select multiple hosts.

**WARNING**

Explicitly assigning a virtual machine to a specific host and disabling migration are mutually exclusive with Red Hat Virtualization high availability.

**IMPORTANT**

If the virtual machine has host devices directly attached to it, and a different host is specified, the host devices from the previous host will be automatically removed from the virtual machine.

5. Select **Allow manual migration only** or **Do not allow migration** from the **Migration Options** drop-down list.
6. Click **OK**.

6.14.8. Manually Migrating Virtual Machines

A running virtual machine can be live migrated to any host within its designated host cluster. Live migration of virtual machines does not cause any service interruption. Migrating virtual machines to a different host is especially useful if the load on a particular host is too high. For live migration prerequisites, see [Live migration prerequisites](#).

For high performance virtual machines and/or virtual machines defined with **Pass-Through Host CPU**, **CPU Pinning**, or **NUMA Pinning**, the default migration mode is **Manual**. Select **Select Host Automatically** so that the virtual machine migrates to the host that offers the best performance.

**NOTE**

When you place a host into maintenance mode, the virtual machines running on that host are automatically migrated to other hosts in the same cluster. You do not need to manually migrate these virtual machines.

**NOTE**

Live migrating virtual machines between different clusters is generally not recommended.

Procedure

1. Click **Compute** → **Virtual Machines** and select a running virtual machine.
2. Click **Migrate**.
3. Use the radio buttons to select whether to **Select Host Automatically** or to **Select Destination Host**, specifying the host using the drop-down list.

**NOTE**

When the **Select Host Automatically** option is selected, the system determines the host to which the virtual machine is migrated according to the load balancing and power management rules set up in the scheduling policy.

4. Click **OK**.

During migration, progress is shown in the **Migration** progress bar. Once migration is complete the **Host** column will update to display the host the virtual machine has been migrated to.

6.14.9. Setting Migration Priority

Red Hat Virtualization Manager queues concurrent requests for migration of virtual machines off of a given host. The load balancing process runs every minute. Hosts already involved in a migration event are not included in the migration cycle until their migration event has completed. When there is a migration request in the queue and available hosts in the cluster to action it, a migration event is triggered in line with the load balancing policy for the cluster.

You can influence the ordering of the migration queue by setting the priority of each virtual machine; for example, setting mission critical virtual machines to migrate before others. Migrations will be ordered by priority; virtual machines with the highest priority will be migrated first.


Setting Migration Priority

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Select the **High Availability** tab.
4. Select **Low**, **Medium**, or **High** from the **Priority** drop-down list.
5. Click **OK**.

6.14.10. Canceling Ongoing Virtual Machine Migrations

A virtual machine migration is taking longer than you expected. You'd like to be sure where all virtual machines are running before you make any changes to your environment.

Procedure

1. Select the migrating virtual machine. It is displayed in **Compute** → **Virtual Machines** with a status of **Migrating from**.
2. Click **More Actions** (), then click **Cancel Migration**.

The virtual machine status returns from **Migrating from** to **Up**.

6.14.11. Event and Log Notification upon Automatic Migration of Highly Available Virtual Servers

When a virtual server is automatically migrated because of the high availability function, the details of an automatic migration are documented in the **Events** tab and in the engine log to aid in troubleshooting, as illustrated in the following examples:

Example 6.4. Notification in the Events Tab of the Administration Portal

Highly Available *Virtual_Machine_Name* failed. It will be restarted automatically.

Virtual_Machine_Name was restarted on Host *Host_Name*

Example 6.5. Notification in the Manager engine.log

This log can be found on the Red Hat Virtualization Manager at `/var/log/ovirt-engine/engine.log`:

Failed to start Highly Available VM. Attempting to restart. VM Name: *Virtual_Machine_Name*, VM Id: *Virtual_Machine_ID_Number*

6.15. IMPROVING UPTIME WITH VIRTUAL MACHINE HIGH AVAILABILITY

6.15.1. What is High Availability?

High availability is recommended for virtual machines running critical workloads. A highly available virtual machine is automatically restarted, either on its original host or another host in the cluster, if its process is interrupted, such as in the following scenarios:

- A host becomes non-operational due to hardware failure.
- A host is put into maintenance mode for scheduled downtime.
- A host becomes unavailable because it has lost communication with an external storage resource.

A highly available virtual machine is not restarted if it is shut down cleanly, such as in the following scenarios:

- The virtual machine is shut down from within the guest.
- The virtual machine is shut down from the Manager.
- The host is shut down by an administrator without being put in maintenance mode first.

With storage domains V4 or later, virtual machines have the additional capability to acquire a lease on a special volume on the storage, enabling a virtual machine to start on another host even if the original host loses power. The functionality also prevents the virtual machine from being started on two different hosts, which may lead to corruption of the virtual machine disks.

With high availability, interruption to service is minimal because virtual machines are restarted within seconds with no user intervention required. High availability keeps your resources balanced by restarting guests on a host with low current resource utilization, or based on any workload balancing or power saving policies that you configure. This ensures that there is sufficient capacity to restart virtual machines at all times.

High Availability and Storage I/O Errors

If a storage I/O error occurs, the virtual machine is paused. You can define how the host handles highly available virtual machines after the connection with the storage domain is reestablished; they can either

be resumed, ungracefully shut down, or remain paused. For more information about these options, see [Virtual Machine High Availability settings explained](#).

6.15.2. High Availability Considerations

A highly available host requires a power management device and fencing parameters. In addition, for a virtual machine to be highly available when its host becomes non-operational, it needs to be started on another available host in the cluster. To enable the migration of highly available virtual machines:


- Power management must be configured for the hosts running the highly available virtual machines.
- The host running the highly available virtual machine must be part of a cluster which has other available hosts.
- The destination host must be running.
- The source and destination host must have access to the data domain on which the virtual machine resides.
- The source and destination host must have access to the same virtual networks and VLANs.
- There must be enough CPUs on the destination host that are not in use to support the virtual machine's requirements.
- There must be enough RAM on the destination host that is not in use to support the virtual machine's requirements.

6.15.3. Configuring a Highly Available Virtual Machine

High availability must be configured individually for each virtual machine.

Procedure

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **High Availability** tab.
4. Select the **Highly Available** check box to enable high availability for the virtual machine.
5. Select the storage domain to hold the virtual machine lease, or select **No VM Lease** to disable the functionality, from the **Target Storage Domain for VM Lease** drop-down list. See [What is high availability](#) for more information about virtual machine leases.



IMPORTANT

This functionality is only available on storage domains that are V4 or later.

6. Select **AUTO_RESUME**, **LEAVE_PAUSED**, or **KILL** from the **Resume Behavior** drop-down list. If you defined a virtual machine lease, **KILL** is the only option available. For more information see [Virtual Machine High Availability settings explained](#).

7. Select **Low**, **Medium**, or **High** from the **Priority** drop-down list. When migration is triggered, a queue is created in which the high priority virtual machines are migrated first. If a cluster is running low on resources, only the high priority virtual machines are migrated.
8. Click **OK**.

6.16. OTHER VIRTUAL MACHINE TASKS

6.16.1. Enabling SAP Monitoring

Enable SAP monitoring on a virtual machine through the Administration Portal.

Enabling SAP Monitoring on Virtual Machines

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click **Edit**.
3. Click the **Custom Properties** tab.
4. Select **sap_agent** from the drop-down list. Ensure the secondary drop-down menu is set to **True**.
If previous properties have been set, select the plus sign to add a new property rule and select **sap_agent**.
5. Click **OK**.

6.16.2. Configuring Red Hat Enterprise Linux 5.4 and later Virtual Machines to use SPICE

SPICE is a remote display protocol designed for virtual environments, which enables you to view a virtualized desktop or server. SPICE delivers a high quality user experience, keeps CPU consumption low, and supports high quality video streaming.

Using SPICE on a Linux machine significantly improves the movement of the mouse cursor on the console of the virtual machine. To use SPICE, the X-Windows system requires additional QXL drivers. The QXL drivers are provided with Red Hat Enterprise Linux 5.4 and later. Earlier versions are not supported. Installing SPICE on a virtual machine running Red Hat Enterprise Linux significantly improves the performance of the graphical user interface.



NOTE

Typically, this is most useful for virtual machines where the user requires the use of the graphical user interface. System administrators who are creating virtual servers may prefer not to configure SPICE if their use of the graphical user interface is minimal.

6.16.2.1. Installing and Configuring QXL Drivers

You must manually install QXL drivers on virtual machines running Red Hat Enterprise Linux 5.4 or later. This is unnecessary for virtual machines running Red Hat Enterprise Linux 6 or Red Hat Enterprise Linux 7 as the QXL drivers are installed by default.

Installing QXL Drivers

1. Log in to a Red Hat Enterprise Linux virtual machine.

2. Install the QXL drivers:

```
# yum install xorg-x11-drv-qxl
```

You can configure QXL drivers using either a graphical interface or the command line. Perform only one of the following procedures.

Configuring QXL drivers in GNOME

1. Click **System**.
2. Click **Administration**.
3. Click **Display**.
4. Click the **Hardware** tab.
5. Click **Video Cards Configure**.
6. Select **qxl** and click **OK**.
7. Restart X-Windows by logging out of the virtual machine and logging back in.

Configuring QXL drivers on the command line

1. Back up **/etc/X11/xorg.conf**:

```
# cp /etc/X11/xorg.conf /etc/X11/xorg.conf.$$backup
```

2. Make the following change to the Device section of **/etc/X11/xorg.conf**:

```
Section "Device"  
Identifier "Videocard0"  
Driver "qxl"  
Endsection
```

6.16.2.2. Configuring a Virtual Machine's Tablet and Mouse to use SPICE

Edit the **/etc/X11/xorg.conf** file to enable SPICE for your virtual machine's tablet devices.

Configuring a Virtual Machine's Tablet and Mouse to use SPICE

1. Verify that the tablet device is available on your guest:

```
# /sbin/lshw -v | grep 'QEMU USB Tablet'
```

If there is no output from the command, do not continue configuring the tablet.

2. Back up **/etc/X11/xorg.conf**:

```
# cp /etc/X11/xorg.conf /etc/X11/xorg.conf.$$backup
```

3. Make the following changes to **/etc/X11/xorg.conf**:

```

Section "ServerLayout"
Identifier "single head configuration"
Screen 0 "Screen0" 0 0
InputDevice "Keyboard0" "CoreKeyboard"
InputDevice "Tablet" "SendCoreEvents"
InputDevice "Mouse" "CorePointer"
EndSection

Section "InputDevice"
Identifier "Mouse"
Driver "void"
#Option "Device" "/dev/input/mice"
#Option "Emulate3Buttons" "yes"
EndSection

Section "InputDevice"
Identifier "Tablet"
Driver "evdev"
Option "Device" "/dev/input/event2"
Option "CorePointer" "true"
EndSection

```

4. Log out and log back into the virtual machine to restart X-Windows.

6.16.3. KVM Virtual Machine Timing Management

Virtualization poses various challenges for virtual machine time keeping. Virtual machines which use the Time Stamp Counter (TSC) as a clock source may suffer timing issues as some CPUs do not have a constant Time Stamp Counter. Virtual machines running without accurate timekeeping can have serious affects on some networked applications as your virtual machine will run faster or slower than the actual time.

KVM works around this issue by providing virtual machines with a paravirtualized clock. The KVM **pvclock** provides a stable source of timing for KVM guests that support it.

Presently, only Red Hat Enterprise Linux 5.4 and later virtual machines fully support the paravirtualized clock.

Virtual machines can have several problems caused by inaccurate clocks and counters:

- Clocks can fall out of synchronization with the actual time which invalidates sessions and affects networks.
- Virtual machines with slower clocks may have issues migrating.

These problems exist on other virtualization platforms and timing should always be tested.



IMPORTANT

The Network Time Protocol (NTP) daemon should be running on the host and the virtual machines. Enable the **ntpd** service and add it to the default startup sequence:

- For Red Hat Enterprise Linux 6

```
# service ntpd start
# chkconfig ntpd on
```

- For Red Hat Enterprise Linux 7

```
# systemctl start ntpd.service
# systemctl enable ntpd.service
```

Using the **ntpd** service should minimize the affects of clock skew in all cases.

The NTP servers you are trying to use must be operational and accessible to your hosts and virtual machines.

Determining if your CPU has the constant Time Stamp Counter

Your CPU has a constant Time Stamp Counter if the **constant_tsc** flag is present. To determine if your CPU has the **constant_tsc** flag run the following command:

```
$ cat /proc/cpuinfo | grep constant_tsc
```

If any output is given your CPU has the **constant_tsc** bit. If no output is given follow the instructions below.

Configuring hosts without a constant Time Stamp Counter

Systems without constant time stamp counters require additional configuration. Power management features interfere with accurate time keeping and must be disabled for virtual machines to accurately keep time with KVM.



IMPORTANT

These instructions are for AMD revision F CPUs only.

If the CPU lacks the **constant_tsc** bit, disable all power management features ([BZ#513138](#)). Each system has several timers it uses to keep time. The TSC is not stable on the host, which is sometimes caused by **cpufreq** changes, deep C state, or migration to a host with a faster TSC. Deep C sleep states can stop the TSC. To prevent the kernel using deep C states append "processor.max_cstate=1" to the kernel boot options in the **grub.conf** file on the host:

```
term Red Hat Enterprise Linux Server (2.6.18-159.el5)
  root (hd0,0)
  kernel /vmlinuz-2.6.18-159.el5 ro root=/dev/VolGroup00/LogVol00 rhgb quiet
processor.max_cstate=1
```

Disable **cpufreq** (only necessary on hosts without the **constant_tsc**) by editing the **/etc/sysconfig/cpuspeed** configuration file and change the **MIN_SPEED** and **MAX_SPEED** variables to the highest frequency available. Valid limits can be found in the

`/sys/devices/system/cpu/cpu/cpufreq/scaling_available_frequencies` files.

Using the `engine-config` tool to receive alerts when hosts drift out of sync.

You can use the `engine-config` tool to configure alerts when your hosts drift out of sync.

There are 2 relevant parameters for time drift on hosts: `EnableHostTimeDrift` and `HostTimeDriftInSec`. `EnableHostTimeDrift`, with a default value of false, can be enabled to receive alert notifications of host time drift. The `HostTimeDriftInSec` parameter is used to set the maximum allowable drift before alerts start being sent.

Alerts are sent once per hour per host.

Using the paravirtualized clock with Red Hat Enterprise Linux virtual machines

For certain Red Hat Enterprise Linux virtual machines, additional kernel parameters are required. These parameters can be set by appending them to the end of the `/kernel` line in the `/boot/grub/grub.conf` file of the virtual machine.



NOTE

The process of configuring kernel parameters can be automated using the `ktune` package

The `ktune` package provides an interactive Bourne shell script, `fix_clock_drift.sh`. When run as the superuser, this script inspects various system parameters to determine if the virtual machine on which it is run is susceptible to clock drift under load. If so, it then creates a new `grub.conf.kvm` file in the `/boot/grub/` directory. This file contains a kernel boot line with additional kernel parameters that allow the kernel to account for and prevent significant clock drift on the KVM virtual machine. After running `fix_clock_drift.sh` as the superuser, and once the script has created the `grub.conf.kvm` file, then the virtual machine's current `grub.conf` file should be backed up manually by the system administrator, the new `grub.conf.kvm` file should be manually inspected to ensure that it is identical to `grub.conf` with the exception of the additional boot line parameters, the `grub.conf.kvm` file should finally be renamed `grub.conf`, and the virtual machine should be rebooted.

The table below lists versions of Red Hat Enterprise Linux and the parameters required for virtual machines on systems without a constant Time Stamp Counter.

Red Hat Enterprise Linux	Additional virtual machine kernel parameters
5.4 AMD64/Intel 64 with the paravirtualized clock	Additional parameters are not required
5.4 AMD64/Intel 64 without the paravirtualized clock	<code>notsc lpj=n</code>
5.4 x86 with the paravirtualized clock	Additional parameters are not required
5.4 x86 without the paravirtualized clock	<code>clocksource=acpi_pm lpj=n</code>
5.3 AMD64/Intel 64	<code>notsc</code>
5.3 x86	<code>clocksource=acpi_pm</code>

Red Hat Enterprise Linux	Additional virtual machine kernel parameters
4.8 AMD64/Intel 64	notsc
4.8 x86	clock=pmtmr
3.9 AMD64/Intel 64	Additional parameters are not required
3.9 x86	Additional parameters are not required

6.16.4. Adding a Trusted Platform Module device

Trusted Platform Module (TPM) devices provide a secure crypto-processor designed to carry out cryptographic operations such as generating cryptographic keys, random numbers, and hashes, or for storing data that can be used to verify software configurations securely. TPM devices are commonly used for disk encryption.

QEMU and libvirt implement support for emulated TPM 2.0 devices, which is what Red Hat Virtualization uses to add TPM devices to Virtual Machines.

Once an emulated TPM device is added to the virtual machine, it can be used as a normal TPM 2.0 device in the guest OS.



IMPORTANT

If there is TPM data stored for the virtual machine and the TPM device is disabled in the virtual machine, the TPM data is permanently removed.


Enabling a TPM device

1. In the **Add Virtual Machine** or **Edit Virtual Machine** screen, click **Show Advanced Options**.
2. In the **Resource Allocation** tab, select the **TPM Device Enabled** check box.

Limitations

The following limitations apply:

- TPM devices can only be used on x86_64 machines with UEFI firmware and PowerPC machines with pSeries firmware installed.
- Virtual machines with TPM devices can not have snapshots with memory.
- While the Manager retrieves and stores TPM data periodically, there is no guarantee that the Manager will always have the latest version of the TPM data.



NOTE

This process can take 120 seconds or more, and you must wait for the process to complete before you can take snapshot of a running virtual machine, clone a running virtual machine, or migrate a running virtual machine.

- TPM devices can only be enabled for virtual machines running RHEL 7 or later and Windows 8.1 or later.
- Virtual machines and templates with TPM data can not be exported or imported.

CHAPTER 7. TEMPLATES

7.1. ABOUT TEMPLATES

A template is a copy of a virtual machine that you can use to simplify the subsequent, repeated creation of similar virtual machines. Templates capture the configuration of software, configuration of hardware, and the software installed on the virtual machine on which the template is based. The virtual machine on which a template is based is known as the source virtual machine.

When you create a template based on a virtual machine, a read-only copy of the virtual machine's disk is created. This read-only disk becomes the base disk image of the new template, and of any virtual machines created based on the template. As such, the template cannot be deleted while any virtual machines created based on the template exist in the environment.

Virtual machines created based on a template use the same NIC type and driver as the original virtual machine, but are assigned separate, unique MAC addresses.

You can create a virtual machine directly from **Compute → Templates**, as well as from **Compute → Virtual Machines**. In **Compute → Templates**, select the required template and click **New VM**. For more information on selecting the settings and controls for the new virtual machine see [Virtual Machine General settings explained](#).

7.2. SEALING VIRTUAL MACHINES IN PREPARATION FOR DEPLOYMENT AS TEMPLATES

This section describes procedures for sealing Linux and Windows virtual machines. Sealing is the process of removing all system-specific details from a virtual machine before creating a template based on that virtual machine. Sealing is necessary to prevent the same details from appearing on multiple virtual machines created based on the same template. It is also necessary to ensure the functionality of other features, such as predictable vNIC order.

7.2.1. Sealing a Linux Virtual Machine for Deployment as a Template

To seal a Linux virtual machine during the template creation process, select the **Seal Template** check box in the **New Template** window. See [Creating a template from an existing virtual machine](#) for details.



NOTE

In RHV 4.4, to seal a RHEL 8 virtual machine for a template, its cluster level must be 4.4 and all hosts in the cluster must be based on RHEL 8. You cannot seal a RHEL 8 virtual machine if you have set its cluster level to 4.3 so it can run on RHEL 7 hosts.

7.2.2. Sealing a Windows Virtual Machine for Deployment as a Template

A template created for Windows virtual machines must be generalized (sealed) before being used to deploy virtual machines. This ensures that machine-specific settings are not reproduced in the template.

Sysprep is used to seal Windows templates before use. **Sysprep** generates a complete unattended installation answer file. Default values for several Windows operating systems are available in the `/usr/share/ovirt-engine/conf/sysprep/` directory. These files act as templates for **Sysprep**. The fields in these files can be copied, pasted, and altered as required. This definition will override any values entered into the **Initial Run** fields of the **Edit Virtual Machine** window.

The Sysprep file can be edited to affect various aspects of the Windows virtual machines created from the template that the Sysprep file is attached to. These include the provisioning of Windows, setting up the required domain membership, configuring the hostname, and setting the security policy.

Replacement strings can be used to substitute values provided in the default files in the `/usr/share/ovirt-engine/conf/sysprep/` directory. For example, "**<Domain><![CDATA[\$JoinDomain\$]></Domain>**" can be used to indicate the domain to join.

7.2.2.1. Prerequisites for Sealing a Windows Virtual Machine



IMPORTANT

Do not reboot the virtual machine while Sysprep is running.

Before starting **Sysprep**, verify that the following settings are configured:

- The Windows virtual machine parameters have been correctly defined.
- If not, click **Edit** in **Compute** → **Virtual Machines** and enter the required information in the **Operating System** and **Cluster** fields.
- The correct product key has been defined in an override file on the Manager.

The override file must be created under `/etc/ovirt-engine/osinfo.conf.d/`, have a filename that puts it after `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties`, and ends in `.properties`. For example, `/etc/ovirt-engine/osinfo.conf.d/10-productkeys.properties`. The last file will have precedence and override any other previous file.

If not, copy the default values for your Windows operating system from `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties` into the override file, and input your values in the `productKey.value` and `sysprepPath.value` fields.

Example 7.1. Windows 7 Default Configuration Values

```
# Windows7(11, OsType.Windows, false),false
os.windows_7.id.value = 11
os.windows_7.name.value = Windows 7
os.windows_7.derivedFrom.value = windows_xp
os.windows_7.sysprepPath.value = ${ENGINE_USR}/conf/sysprep/sysprep.w7
os.windows_7.productKey.value =
os.windows_7.devices.audio.value = ich6
os.windows_7.devices.diskInterfaces.value.3.3 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.devices.diskInterfaces.value.3.4 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.devices.diskInterfaces.value.3.5 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.isTimezoneTypeInteger.value = false
```

7.2.2.2. Sealing a Windows 7, Windows 2008, or Windows 2012 Virtual Machine for Deployment as Template

Seal a Windows 7, Windows 2008, or Windows 2012 virtual machine before creating a template to use to deploy virtual machines.


Procedure

1. On the Windows virtual machine, launch **Sysprep** from **C:\Windows\System32\sysprep\sysprep.exe**.
2. Enter the following information into **Sysprep**:
 - Under **System Cleanup Action**, select **Enter System Out-of-Box-Experience (OOBE)**.
 - Select the **Generalize** check box if you need to change the computer's system identification number (SID).
 - Under **Shutdown Options**, select **Shutdown**.
3. Click **OK** to complete the sealing process; the virtual machine shuts down automatically upon completion.

The Windows 7, Windows 2008, or Windows 2012 virtual machine is sealed and ready to create a template to use for deploying virtual machines.

7.3. CREATING A TEMPLATE

Create a template from an existing virtual machine to use as a blueprint for creating additional virtual machines.




NOTE

In RHV 4.4, to seal a RHEL 8 virtual machine for a template, its cluster level must be 4.4 and all hosts in the cluster must be based on RHEL 8. You cannot seal a RHEL 8 virtual machine if you have set its cluster level to 4.3 so it can run on RHEL 7 hosts.


When you create a template, you specify the format of the disk to be raw or QCOW2:

- QCOW2 disks are thin provisioned.
- Raw disks on file storage are thin provisioned.
- Raw disks on block storage are preallocated.

Creating a Template

1. Click **Compute** → **Virtual Machines** and select the source virtual machine.
2. Ensure the virtual machine is powered down and has a status of **Down**.
3. Click **More Actions** (), then click **Make Template**. For more details on all fields in the **New Template** window, see [Explanation of Settings in the New Template and Edit Template Windows](#).
4. Enter a **Name**, **Description**, and **Comment** for the template.
5. Select the cluster with which to associate the template from the **Cluster** drop-down list. By default, this is the same as that of the source virtual machine.
6. Optionally, select a CPU profile for the template from the **CPU Profile** drop-down list.

7. Optionally, select the **Create as a Template Sub-Version** check box, select a **Root Template**, and enter a **Sub-Version Name** to create the new template as a sub-template of an existing template.
8. In the **Disks Allocation** section, enter an alias for the disk in the **Alias** text field. Select the disk format in the **Format** drop-down, the storage domain on which to store the disk from the **Target** drop-down, and the disk profile in the **Disk Profile** drop-down. By default, these are the same as those of the source virtual machine.
9. Select the **Allow all users to access this Template** check box to make the template public.
10. Select the **Copy VM permissions** check box to copy the permissions of the source virtual machine to the template.
11. Select the **Seal Template** check box (Linux only) to seal the template.



NOTE

Sealing, which uses the **virt-sysprep** command, removes system-specific details from a virtual machine before creating a template based on that virtual machine. This prevents the original virtual machine's details from appearing in subsequent virtual machines that are created using the same template. It also ensures the functionality of other features, such as predictable vNIC order. See [virt-sysprep operations](#) for more information.

12. Click **OK**.

The virtual machine displays a status of **Image Locked** while the template is being created. The process of creating a template may take up to an hour depending on the size of the virtual disk and the capabilities of your storage hardware. When complete, the template is added to the **Templates** tab. You can now create new virtual machines based on the template.



NOTE

When a template is made, the virtual machine is copied so that both the existing virtual machine and its template are usable after template creation.

7.4. EDITING A TEMPLATE

Once a template has been created, its properties can be edited. Because a template is a copy of a virtual machine, the options available when editing a template are identical to those in the **Edit Virtual Machine** window.

Procedure

1. Click **Compute** → **Templates** and select a template.
2. Click **Edit**.
3. Change the necessary properties. Click **Show Advanced Options** and edit the template's settings as required. The settings that appear in the **Edit Template** window are identical to those in the **Edit Virtual Machine** window, but with the relevant fields only. See [Explanation of Settings in the New Virtual Machine and Edit Virtual Machine Windows](#) for details.
4. Click **OK**.

7.5. DELETING A TEMPLATE

If you have used a template to create a virtual machine using the thin provisioning storage allocation option, the template cannot be deleted as the virtual machine needs it to continue running. However, cloned virtual machines do not depend on the template they were cloned from and the template can be deleted.

Deleting a Template

1. Click **Compute** → **Templates** and select a template.
2. Click **Remove**.
3. Click **OK**.

7.6. EXPORTING TEMPLATES

7.6.1. Migrating Templates to the Export Domain



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center. See the [Importing Existing Storage Domains](#) section in the *Red Hat Virtualization Administration Guide* for information on importing storage domains.

Export templates into the export domain to move them to another data domain, either in the same Red Hat Virtualization environment, or another one. This procedure requires access to the Administration Portal.

Exporting Individual Templates to the Export Domain

1. Click **Compute** → **Templates** and select a template.
2. Click **Export**.
3. Select the **Force Override** check box to replace any earlier version of the template on the export domain.
4. Click **OK** to begin exporting the template; this may take up to an hour, depending on the virtual disk size and your storage hardware.

Repeat these steps until the export domain contains all the templates to migrate before you start the import process.

1. Click **Storage** → **Domains** and select the export domain.
2. Click the domain name to see the details view.
3. Click the **Template Import** tab to view all exported templates in the export domain.

7.6.2. Copying a Template's Virtual Hard Disk

If you are moving a virtual machine that was created from a template with the thin provisioning storage allocation option selected, the template's disks must be copied to the same storage domain as that of the virtual disk. This procedure requires access to the Administration Portal.

Copying a Virtual Hard Disk

1. Click **Storage** → **Disks**.
2. Select the template disk(s) to copy.
3. Click **Copy**.
4. Select the **Target** data domain from the drop-down list(s).
5. Click **OK**.

A copy of the template's virtual hard disk has been created, either on the same, or a different, storage domain. If you were copying a template disk in preparation for moving a virtual hard disk, you can now move the virtual hard disk.

7.7. IMPORTING TEMPLATES

7.7.1. Importing a Template into a Data Center



NOTE

The export storage domain is deprecated. Storage data domains can be unattached from a data center and imported to another data center in the same environment, or in a different environment. Virtual machines, floating virtual disks, and templates can then be uploaded from the imported storage domain to the attached data center. See the [Importing Existing Storage Domains](#) section in the *Red Hat Virtualization Administration Guide* for information on importing storage domains.

Import templates from a newly attached export domain. This procedure requires access to the Administration Portal.

Importing a Template into a Data Center

1. Click **Storage** → **Domains** and select the newly attached export domain.
2. Click the domain name to go to the details view.
3. Click the **Template Import** tab and select a template.
4. Click **Import**.
5. Use the drop-down lists to select the **Target Cluster** and **CPU Profile**.
6. Select the template to view its details, then click the **Disks** tab and select the **Storage Domain** to import the template into.
7. Click **OK**.
8. If the **Import Template Conflict** window appears, enter a **New Name** for the template, or select the **Apply to all** check box and enter a **Suffix to add to the cloned Templates**. Click **OK**.

9. Click **Close**.

The template is imported into the destination data center. This can take up to an hour, depending on your storage hardware. You can view the import progress in the **Events** tab.

Once the importing process is complete, the templates will be visible in **Compute → Templates**. The templates can create new virtual machines, or run existing imported virtual machines based on that template.

7.7.2. Importing a Virtual Disk from an OpenStack Image Service as a Template

Virtual disks managed by an OpenStack Image Service can be imported into the Red Hat Virtualization Manager if that OpenStack Image Service has been added to the Manager as an external provider. This procedure requires access to the Administration Portal.

1. Click **Storage → Domains** and select the OpenStack Image Service domain.
2. Click the storage domain name to go to the details view.
3. Click the **Images** tab and select the image to import.
4. Click **Import**.



NOTE

If you are importing an image from a Glance storage domain, you have the option of specifying the template name. OpenStack Glance is now deprecated. This functionality will be removed in a later release.

5. Select the **Data Center** into which the virtual disk will be imported.
6. Select the storage domain in which the virtual disk will be stored from the **Domain Name** drop-down list.
7. Optionally, select a **Quota** to apply to the virtual disk.
8. Select the **Import as Template** check box.
9. Select the **Cluster** in which the virtual disk will be made available as a template.
10. Click **OK**.

The image is imported as a template and is displayed in the **Templates** tab. You can now create virtual machines based on the template.

7.8. TEMPLATES AND PERMISSIONS

7.8.1. Managing System Permissions for a Template

As the **SuperUser**, the system administrator manages all aspects of the Administration Portal. More specific administrative roles can be assigned to other users. These restricted administrator roles are useful for granting a user administrative privileges that limit them to a specific resource. For example, a **DataCenterAdmin** role has administrator privileges only for the assigned data center with the exception of the storage for that data center, and a **ClusterAdmin** has administrator privileges only for the assigned cluster.

A template administrator is a system administration role for templates in a data center. This role can be applied to specific virtual machines, to a data center, or to the whole virtualized environment; this is useful to allow different users to manage certain virtual resources.

The template administrator role permits the following actions:

- Create, edit, export, and remove associated templates.
- Import and export templates.



NOTE

You can only assign roles and permissions to existing users.

7.8.2. Template Administrator Roles Explained

The table below describes the administrator roles and privileges applicable to template administration.

Table 7.1. Red Hat Virtualization System Administrator Roles

Role	Privileges	Notes
TemplateAdmin	Can perform all operations on templates.	Has privileges to create, delete and configure a template's storage domain and network details, and to move templates between domains.
NetworkAdmin	Network Administrator	Can configure and manage networks attached to templates.

7.8.3. Assigning an Administrator or User Role to a Resource

Assign administrator or user roles to resources to allow users to access or manage that resource.

Procedure

1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the resource's name to go to the details view.
3. Click the **Permissions** tab to list the assigned users, the user's role, and the inherited permissions for the selected resource.
4. Click **Add**.
5. Enter the name or user name of an existing user into the **Search** text box and click **Go**. Select a user from the resulting list of possible matches.
6. Select a role from the **Role to Assign:** drop-down list.
7. Click **OK**.

You have assigned a role to a user; the user now has the inherited permissions of that role enabled for that resource.

7.8.4. Removing an Administrator or User Role from a Resource

Remove an administrator or user role from a resource; the user loses the inherited permissions associated with the role for that resource.

Removing a Role from a Resource


1. Use the resource tabs, tree mode, or the search function to find and select the resource in the results list.
2. Click the resource's name to go to the details view.
3. Click the **Permissions** tab to list the assigned users, the user's role, and the inherited permissions for the selected resource.
4. Select the user to remove from the resource.
5. Click **Remove**. The **Remove Permission** window opens to confirm permissions removal.
6. Click **OK**.

You have removed the user's role, and the associated permissions, from the resource.

7.9. USING CLOUD-INIT TO AUTOMATE THE CONFIGURATION OF VIRTUAL MACHINES

Cloud-Init is a tool for automating the initial setup of virtual machines such as configuring the host name, network interfaces, and authorized keys. It can be used when provisioning virtual machines that have been deployed based on a template to avoid conflicts on the network.

To use this tool, the **cloud-init** package must first be installed on the virtual machine. Once installed, the Cloud-Init service starts during the boot process to search for instructions on what to configure. You can then use options in the **Run Once** window to provide these instructions one time only, or options in the **New Virtual Machine**, **Edit Virtual Machine** and **Edit Template** windows to provide these instructions every time the virtual machine starts.



NOTE

Alternatively, you can configure Cloud-Init with [Ansible](#), [Python](#), [Java](#), or [Ruby](#).

7.9.1. Cloud-Init Use Case Scenarios

Cloud-Init can be used to automate the configuration of virtual machines in a variety of scenarios. Several common scenarios are as follows:

- **Virtual Machines Created Based on Templates**
You can use the Cloud-Init options in the **Initial Run** section of the **Run Once** window to initialize a virtual machine that was created based on a template. This allows you to customize the virtual machine the first time that virtual machine is started.
- **Virtual Machine Templates**

You can use the **Use Cloud-Init/Sysprep** options in the **Initial Run** tab of the **Edit Template** window to specify options for customizing virtual machines created based on that template.

- **Virtual Machine Pools**

You can use the **Use Cloud-Init/Sysprep** options in the **Initial Run** tab of the **New Pool** window to specify options for customizing virtual machines taken from that virtual machine pool. This allows you to specify a set of standard settings that will be applied every time a virtual machine is taken from that virtual machine pool. You can inherit or override the options specified for the template on which the virtual machine is based, or specify options for the virtual machine pool itself.

7.9.2. Installing Cloud-Init

This procedure describes how to install Cloud-Init on a virtual machine. Once Cloud-Init is installed, you can create a template based on this virtual machine. Virtual machines created based on this template can leverage Cloud-Init functions, such as configuring the host name, time zone, root password, authorized keys, network interfaces, DNS service, etc on boot.

Installing Cloud-Init

1. Log in to the virtual machine.
2. Enable the repositories:
 - For Red Hat Enterprise Linux 6:

```
# subscription-manager repos \
  --enable=rhel-6-server-rpms \
  --enable=rhel-6-server-rh-common-rpms
```

- For Red Hat Enterprise Linux 7:

```
# subscription-manager repos \
  --enable=rhel-7-server-rpms \
  --enable=rhel-7-server-rh-common-rpms
```

- For Red Hat Enterprise Linux 8, you normally do not need to enable repositories to install Cloud-Init. The Cloud-Init package is part of the AppStream repo, **rhel-8-for-x86_64-appstream-rpms**, which is enabled by default in Red Hat Enterprise Linux 8.

3. Install the **cloud-init** package and dependencies:

```
# dnf install cloud-init
```



NOTE

For versions of Red Hat Enterprise Linux earlier than version 8, use the command **yum install cloud-init** instead of **dnf install cloud-init**.

7.9.3. Using Cloud-Init to Prepare a Template

As long as the **cloud-init** package is installed on a Linux virtual machine, you can use the virtual machine to make a cloud-init enabled template. Specify a set of standard settings to be included in a template as described in the following procedure or, alternatively, skip the Cloud-Init settings steps and configure

them when creating a virtual machine based on this template.



NOTE

While the following procedure outlines how to use Cloud-Init when preparing a template, the same settings are also available in the **New Virtual Machine**, **Edit Template**, and **Run Once** windows.

Using Cloud-Init to Prepare a Template

1. Click **Compute** → **Templates** and select a template.
2. Click **Edit**.
3. Click **Show Advanced Options**
4. Click the **Initial Run** tab and select the **Use Cloud-Init/Sysprep** check box.
5. Enter a host name in the **VM Hostname** text field.
6. Select the **Configure Time Zone** check box and select a time zone from the **Time Zone** drop-down list.
7. Expand the **Authentication** section.
 - Select the **Use already configured password** check box to use the existing credentials, or clear that check box and enter a root password in the **Password** and **Verify Password** text fields to specify a new root password.
 - Enter any SSH keys to be added to the authorized hosts file on the virtual machine in the **SSH Authorized Keys** text area.
 - Select the **Regenerate SSH Keys** check box to regenerate SSH keys for the virtual machine.
8. Expand the **Networks** section.
 - Enter any DNS servers in the **DNS Servers** text field.
 - Enter any DNS search domains in the **DNS Search Domains** text field.
 - Select the **In-guest Network Interface** check box and use the **+ Add new** and **- Remove selected** buttons to add or remove network interfaces to or from the virtual machine.



IMPORTANT

You must specify the correct network interface name and number (for example, **eth0**, **eno3**, **enp0s**). Otherwise, the virtual machine's interface connection will be up, but it will not have the **cloud-init** network configuration.

9. Expand the **Custom Script** section and enter any custom scripts in the **Custom Script** text area.
10. Click **OK**.

You can now provision new virtual machines using this template.

7.9.4. Using Cloud-Init to Initialize a Virtual Machine

Use Cloud-Init to automate the initial configuration of a Linux virtual machine. You can use the Cloud-Init fields to configure a virtual machine's host name, time zone, root password, authorized keys, network interfaces, and DNS service. You can also specify a custom script, a script in YAML format, to run on boot. The custom script allows for additional Cloud-Init configuration that is supported by Cloud-Init but not available in the Cloud-Init fields. For more information on custom script examples, see [Cloud config examples](#).

Using Cloud-Init to Initialize a Virtual Machine

This procedure starts a virtual machine with a set of Cloud-Init settings. If the relevant settings are included in the template the virtual machine is based on, review the settings, make changes where appropriate, and click **OK** to start the virtual machine.

1. Click **Compute** → **Virtual Machines** and select a virtual machine.
2. Click the **Run** drop-down button and select **Run Once**.
3. Expand the **Initial Run** section and select the **Cloud-Init** check box.
4. Enter a host name in the **VM Hostname** text field.
5. Select the **Configure Time Zone** check box and select a time zone from the **Time Zone** drop-down menu.
6. Select the **Use already configured password** check box to use the existing credentials, or clear that check box and enter a root password in the **Password** and **Verify Password** text fields to specify a new root password.
7. Enter any SSH keys to be added to the authorized hosts file on the virtual machine in the **SSH Authorized Keys** text area.
8. Select the **Regenerate SSH Keys** check box to regenerate SSH keys for the virtual machine.
9. Enter any DNS servers in the **DNS Servers** text field.
10. Enter any DNS search domains in the **DNS Search Domains** text field.
11. Select the **Network** check box and use the + and - buttons to add or remove network interfaces to or from the virtual machine.



IMPORTANT

You must specify the correct network interface name and number (for example, **eth0**, **eno3**, **enp0s**). Otherwise, the virtual machine's interface connection will be up, but the **cloud-init** network configuration will not be defined in it.

12. Enter a custom script in the **Custom Script** text area. Make sure the values specified in the script are appropriate. Otherwise, the action will fail.
13. Click **OK**.

**NOTE**

To check if a virtual machine has Cloud-Init installed, select a virtual machine and click the **Applications** sub-tab. Only shown if the guest agent is installed.

7.10. USING SYSPREP TO AUTOMATE THE CONFIGURATION OF VIRTUAL MACHINES

Sysprep is a tool used to automate the setup of Windows virtual machines, for example, configuring host names, network interfaces, authorized keys, set up users, or to connect to Active Directory. **Sysprep** is installed with every version of Windows.

Red Hat Virtualization enhances **Sysprep** by exploiting virtualization technology to deploy virtual workstations based on a single template. Red Hat Virtualization builds a tailored auto-answer file for each virtual workstation.

Sysprep generates a complete unattended installation answer file. Default values for several Windows operating systems are available in the `/usr/share/ovirt-engine/conf/sysprep/` directory. You can also create a custom **Sysprep** file and reference it from the `osinfo` file in the `/etc/ovirt-engine/osinfo.conf.d/` directory. These files act as templates for **Sysprep**. The fields in these files can be copied and edited as required. This definition will override any values entered into the **Initial Run** fields of the **Edit Virtual Machine** window.

You can create a custom **sysprep** file when creating a pool of Windows virtual machines, to accommodate various operating systems and domains. See [Creating a Virtual Machine Pool](#) in the *Administration Guide* for details.

The override file must be created under `/etc/ovirt-engine/osinfo.conf.d/`, have a filename that puts it after `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties`, and ends in `.properties`. For example, `/etc/ovirt-engine/osinfo.conf.d/10-productkeys.properties`. The last file will have precedence and override any other previous file.

Copy the default values for your Windows operating system from `/etc/ovirt-engine/osinfo.conf.d/00-defaults.properties` into the override file, and input your values in the `productKey.value` and `sysprepPath.value` fields.

Example 7.2. Windows 7 Default Configuration Values

```
# Windows7(11, OsType.Windows, false),false
os.windows_7.id.value = 11
os.windows_7.name.value = Windows 7
os.windows_7.derivedFrom.value = windows_xp
os.windows_7.sysprepPath.value = ${ENGINE_USR}/conf/sysprep/sysprep.w7
os.windows_7.productKey.value =
os.windows_7.devices.audio.value = ich6
os.windows_7.devices.diskInterfaces.value.3.3 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.devices.diskInterfaces.value.3.4 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.devices.diskInterfaces.value.3.5 = IDE, VirtIO_SCSI, VirtIO
os.windows_7.isTimezoneTypeInteger.value = false
```

7.10.1. Configuring Sysprep on a Template

You can use this procedure to specify a set of standard **Sysprep** settings to include in the template, alternatively you can configure the **Sysprep** settings when creating a virtual machine based on this template.

Replacement strings can be used to substitute values provided in the default files in the `/usr/share/ovirt-engine/conf/sysprep/` directory. For example, "`<Domain><![CDATA[$JoinDomain$]></Domain>`" can be used to indicate the domain to join.



IMPORTANT

Do not reboot the virtual machine while **Sysprep** is running.

Prerequisites

- The Windows virtual machine parameters have been correctly defined. If not, click **Compute** → **Virtual Machines**, click **Edit**, and enter the required information in the **Operating System** and **Cluster** fields.
- The correct product key has been defined in an override file on the Manager.

Using Sysprep to Prepare a Template

1. Build the Windows virtual machine with the required patches and software.
2. Seal the Windows virtual machine. See [Sealing Virtual Machines in Preparation for Deployment as Templates](#)
3. Create a template based on the Windows virtual machine. See [Creating a template from an existing virtual machine](#)
4. Update the **Sysprep** file with a text editor if additional changes are required.

You can now provision new virtual machines using this template.

7.10.2. Using Sysprep to Initialize a Virtual Machine

Use **Sysprep** to automate the initial configuration of a Windows virtual machine. You can use the **Sysprep** fields to configure a virtual machine's host name, time zone, root password, authorized keys, network interfaces, and DNS service.

Using Sysprep to Initialize a Virtual Machine

This procedure starts a virtual machine with a set of **Sysprep** settings. If the relevant settings are included in the template the virtual machine is based on, review the settings and make changes where required.

1. Create a new Windows virtual machine based on a template of the required Windows virtual machine. See [Creating a Virtual Machine Based on a Template](#) .
2. Click **Compute** → **Virtual Machines** and select the virtual machine.
3. Click the **Run** drop-down button and select **Run Once**.
4. Expand the **Boot Options** section, select the **Attach Floppy** check box, and select the **[sysprep]** option.
5. Select the **Attach CD** check box and select the required Windows ISO from the drop-down list.

6. Move the **CD-ROM** to the top of the **Boot Sequence** field.
7. Configure any further **Run Once** options as required. See [Virtual Machine Run Once settings explained](#) for more details.
8. Click **OK**.

7.11. CREATING A VIRTUAL MACHINE BASED ON A TEMPLATE

Create a virtual machine from a template to enable the virtual machines to be pre-configured with an operating system, network interfaces, applications and other resources.



NOTE

Virtual machines created from a template depend on that template. So you cannot remove a template from the Manager if a virtual machine was created from that template. However, you can clone a virtual machine from a template to remove the dependency on that template.



NOTE

If the BIOS type of the virtual machine differs from the BIOS type of the template, the Manager might change devices in the virtual machine, possibly preventing the operating system from booting. For example, if the template uses IDE disks and the i440fx chipset, changing the BIOS type to the Q35 chipset automatically changes the IDE disks to SATA disks. So configure the chipset and BIOS type to match the chipset and BIOS type of the template.

Creating a Virtual Machine Based on a Template

1. Click **Compute** → **Virtual Machines**.
2. Click **New**.
3. Select the **Cluster** on which the virtual machine will run.
4. Select a template from the **Template** list.
5. Enter a **Name**, **Description**, and any **Comments**, and accept the default values inherited from the template in the rest of the fields. You can change them if needed.
6. Click the **Resource Allocation** tab.
7. Select the **Thin** or **Clone** radio button in the **Storage Allocation** area. If you select **Thin**, the disk format is QCOW2. If you select **Clone**, select either **QCOW2** or **Raw** for disk format.
8. Use the **Target** drop-down list to select the storage domain on which the virtual machine's virtual disk will be stored.
9. Click **OK**.

The virtual machine is displayed in the **Virtual Machines** tab.

Additional Resources

- [Creating a cloned virtual machine based on a template](#)
- [UEFI and the Q35 chipset](#) in the *Administration Guide*.

7.12. CREATING A CLONED VIRTUAL MACHINE BASED ON A TEMPLATE

Cloned virtual machines are based on templates and inherit the settings of the template. A cloned virtual machine does not depend on the template on which it was based after it has been created. This means the template can be deleted if no other dependencies exist.



NOTE

If you clone a virtual machine from a template, the name of the template on which that virtual machine was based is displayed in the **General** tab of the **Edit Virtual Machine** window for that virtual machine. If you change the name of that template, the name of the template in the **General** tab will also be updated. However, if you delete the template from the Manager, the original name of that template will be displayed instead.

Cloning a Virtual Machine Based on a Template

1. Click **Compute** → **Virtual Machines**.
2. Click **New**.
3. Select the **Cluster** on which the virtual machine will run.
4. Select a template from the **Based on Template** drop-down menu.
5. Enter a **Name**, **Description** and any **Comments**. You can accept the default values inherited from the template in the rest of the fields, or change them if required.
6. Click the **Resource Allocation** tab.
7. Select the **Clone** radio button in the **Storage Allocation** area.
8. Select the disk format from the **Format** drop-down list. This affects the speed of the clone operation and the amount of disk space the new virtual machine initially requires.
 - **QCOW2** (Default)
 - Faster clone operation
 - Optimized use of storage capacity
 - Disk space allocated only as required
 - **Raw**
 - Slower clone operation
 - Optimized virtual machine read and write operations
 - All disk space requested in the template is allocated at the time of the clone operation

9. Use the **Target** drop-down menu to select the storage domain on which the virtual machine's virtual disk will be stored.
10. Click **OK**.

**NOTE**

Cloning a virtual machine may take some time. A new copy of the template's disk must be created. During this time, the virtual machine's status is first **Image Locked**, then **Down**.

The virtual machine is created and displayed in the **Virtual Machines** tab. You can now assign users to it, and can begin using it when the clone operation is complete.

APPENDIX A. REFERENCE: SETTINGS IN ADMINISTRATION PORTAL AND VM PORTAL WINDOWS

A.1. EXPLANATION OF SETTINGS IN THE NEW VIRTUAL MACHINE AND EDIT VIRTUAL MACHINE WINDOWS

A.1.1. Virtual Machine General Settings Explained

The following table details the options available on the **General** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.1. Virtual Machine: General Settings

Field Name	Description	Power cycle required?
Cluster	The name of the host cluster to which the virtual machine is attached. Virtual machines are hosted on any physical machine in that cluster in accordance with policy rules.	Yes. Cross-cluster migration is for emergency use only. Moving clusters requires the virtual machine to be down.
Template	<p>The template on which the virtual machine is based. This field is set to Blank by default, which allows you to create a virtual machine on which an operating system has not yet been installed. Templates are displayed as Name Sub-version name (Sub-version number). Each new version is displayed with a number in brackets that indicates the relative order of the version, with a higher number indicating a more recent version.</p> <p>The version name is displayed as base version if it is the root template of the template version chain.</p> <p>When the virtual machine is stateless, there is an option to select the latest version of the template. This option means that anytime a new version of this template is created, the virtual machine is automatically recreated on restart based on the latest template.</p>	Not applicable. This setting is for provisioning new virtual machines only.
Operating System	The operating system. Valid values include a range of Red Hat Enterprise Linux and Windows variants.	Yes. Potentially changes the virtual hardware.

Field Name	Description	Power cycle required?
Instance Type	<p>The instance type on which the virtual machine's hardware configuration can be based. This field is set to Custom by default, which means the virtual machine is not connected to an instance type. The other options available from this drop down menu are Large, Medium, Small, Tiny, XLarge, and any custom instance types that the Administrator has created.</p> <p>Other settings that have a chain link icon next to them are pre-filled by the selected instance type. If one of these values is changed, the virtual machine will be detached from the instance type and the chain icon will appear broken. However, if the changed setting is restored to its original value, the virtual machine will be reattached to the instance type and the links in the chain icon will rejoin.</p> <p>NOTE: Support for instance types is now deprecated, and will be removed in a future release.</p>	Yes.
Optimized for	<p>The type of system for which the virtual machine is to be optimized. There are three options: Server, Desktop, and High Performance; by default, the field is set to Server. Virtual machines optimized to act as servers have no sound card, use a cloned disk image, and are not stateless. Virtual machines optimized to act as desktop machines do have a sound card, use an image (thin allocation), and are stateless. Virtual machines optimized for high performance have a number of configuration changes. See Configuring High Performance Virtual Machines Templates and Pools.</p>	Yes.
Name	<p>The name of the virtual machine. The name must be a unique name within the data center and must not contain any spaces, and must contain at least one character from A-Z or 0-9. The maximum length of a virtual machine name is 255 characters. The name can be reused in different data centers in the environment.</p>	Yes.
VM ID	<p>The virtual machine ID. The virtual machine's creator can set a custom ID for that virtual machine. The custom ID must contain only numbers, in the format, 00000000-0000-0000-0000-00000000.</p> <p>If no ID is specified during creation a UUID will be automatically assigned. For both custom and automatically-generated IDs, changes are not possible after virtual machine creation.</p>	Yes.

Field Name	Description	Power cycle required?
Description	A meaningful description of the new virtual machine.	No.
Comment	A field for adding plain text human-readable comments regarding the virtual machine.	No.
Affinity Labels	Add or remove a selected Affinity Label .	No.
Stateless	Select this check box to run the virtual machine in stateless mode. This mode is used primarily for desktop virtual machines. Running a stateless desktop or server creates a new COW layer on the virtual machine hard disk image where new and changed data is stored. Shutting down the stateless virtual machine deletes the new COW layer which includes all data and configuration changes, and returns the virtual machine to its original state. Stateless virtual machines are useful when creating machines that need to be used for a short time, or by temporary staff.	Not applicable.
Start in Pause Mode	Select this check box to always start the virtual machine in pause mode. This option is suitable for virtual machines which require a long time to establish a SPICE connection; for example, virtual machines in remote locations.	Not applicable.
Delete Protection	Select this check box to make it impossible to delete the virtual machine. It is only possible to delete the virtual machine if this check box is not selected.	No.
Sealed	Select this check box to seal the created virtual machine. This option eliminates machine-specific settings from virtual machines that are provisioned from the template. For more information about the sealing process, see Sealing a Windows Virtual Machine for Deployment as a Template	No.
Instance Images	Click Attach to attach a floating disk to the virtual machine, or click Create to add a new virtual disk. Use the plus and minus buttons to add or remove additional virtual disks. Click Edit to change the configuration of a virtual disk that has already been attached or created.	No.
Instantiate VM network interfaces by picking a vNIC profile.	Add a network interface to the virtual machine by selecting a vNIC profile from the nic1 drop-down list. Use the plus and minus buttons to add or remove additional network interfaces.	No.

A.1.2. Virtual Machine System Settings Explained

CPU Considerations

- **For non-CPU-intensive workloads**, you can run virtual machines with a total number of processor cores greater than the number of cores in the host (the number of processor cores for a single virtual machine must not exceed the number of cores in the host). The following benefits can be achieved:
 - You can run a greater number of virtual machines, which reduces hardware requirements.
 - You can configure virtual machines with CPU topologies that are otherwise not possible, such as when the number of virtual cores is between the number of host cores and the number of host threads.
- **For best performance, and especially for CPU-intensive workloads**, you should use the same topology in the virtual machine as in the host, so the host and the virtual machine expect the same cache usage. When the host has hyperthreading enabled, QEMU treats the host's hyperthreads as cores, so the virtual machine is not aware that it is running on a single core with multiple threads. This behavior might impact the performance of a virtual machine, because a virtual core that actually corresponds to a hyperthread in the host core might share a single cache with another hyperthread in the same host core, while the virtual machine treats it as a separate core.

The following table details the options available on the **System** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.2. Virtual Machine: System Settings

Field Name	Description	Power cycle required?
Memory Size	The amount of memory assigned to the virtual machine. When allocating memory, consider the processing and storage needs of the applications that are intended to run on the virtual machine.	If OS supports hotplugging, no. Otherwise, yes.
Maximum Memory	The maximum amount of memory that can be assigned to the virtual machine. Maximum guest memory is also constrained by the selected guest architecture and the cluster compatibility level.	If OS supports hotplugging, no. Otherwise, yes.
Total Virtual CPUs	The processing power allocated to the virtual machine as CPU Cores. For high performance, do not assign more cores to a virtual machine than are present on the physical host.	If OS supports hotplugging, no. Otherwise, yes.
Virtual Sockets	The number of CPU sockets for the virtual machine. Do not assign more sockets to a virtual machine than are present on the physical host.	If OS supports hotplugging, no. Otherwise, yes.
Cores per Virtual Socket	The number of cores assigned to each virtual socket.	If OS supports hotplugging, no. Otherwise, yes.

Field Name	Description	Power cycle required?
Threads per Core	The number of threads assigned to each core. Increasing the value enables simultaneous multi-threading (SMT). IBM POWER8 supports up to 8 threads per core. For x86 and x86_64 (Intel and AMD) CPU types, the recommended value is 1, unless you want to replicate the exact host topology, which you can do using CPU pinning. For more information, see Pinning CPU .	If OS supports hotplugging, no. Otherwise, yes.
Chipset/Firmware Type	Specifies the chipset and firmware type. Defaults to the cluster's default chipset and firmware type. Options are: <ul style="list-style-type: none"> ● I440FX Chipset with BIOS Legacy BIOS ● Q35 Chipset with BIOS BIOS without UEFI (Default for clusters with compatibility version 4.4) ● Q35 Chipset with UEFI BIOS with UEFI (Default for clusters with compatibility version 4.7) ● Q35 Chipset with UEFI SecureBoot UEFI with SecureBoot, which authenticates the digital signatures of the boot loader For more information, see UEFI and the Q35 chipset in the <i>Administration Guide</i> .	Yes.
Custom Emulated Machine	This option allows you to specify the machine type. If changed, the virtual machine will only run on hosts that support this machine type. Defaults to the cluster's default machine type.	Yes.
Custom CPU Type	This option allows you to specify a CPU type. If changed, the virtual machine will only run on hosts that support this CPU type. Defaults to the cluster's default CPU type.	Yes.
Hardware Clock Time Offset	This option sets the time zone offset of the guest hardware clock. For Windows, this should correspond to the time zone set in the guest. Most default Linux installations expect the hardware clock to be GMT+00:00.	Yes.

Field Name	Description	Power cycle required?
Custom Compatibility Version	The compatibility version determines which features are supported by the cluster, as well as, the values of some properties and the emulated machine type. By default, the virtual machine is configured to run in the same compatibility mode as the cluster as the default is inherited from the cluster. In some situations the default compatibility mode needs to be changed. An example of this is if the cluster has been updated to a later compatibility version but the virtual machines have not been restarted. These virtual machines can be set to use a custom compatibility mode that is older than that of the cluster. See Changing the Cluster Compatibility Version in the <i>Administration Guide</i> for more information.	Yes.
Serial Number Policy	Override the system-level and cluster-level policies for assigning a serial numbers to virtual machines. Apply a policy that is unique to this virtual machine: <ul style="list-style-type: none"> ● System Default: Use the system-wide defaults, which are configured in the Manager database using the engine configuration tool and the DefaultSerialNumberPolicy and DefaultCustomSerialNumber key names. The default value for DefaultSerialNumberPolicy is to use the Host ID. See Scheduling Policies in the <i>Administration Guide</i> for more information. ● Host ID: Set this virtual machine's serial number to the UUID of the host. ● Vm ID: Set this virtual machine's serial number to the UUID of this virtual machine. ● Custom serial number: Set this virtual machine's serial number to the value you specify in the following Custom Serial Number parameter. 	Yes.
Custom Serial Number	Specify the custom serial number to apply to this virtual machine.	Yes.

A.1.3. Virtual Machine Initial Run Settings Explained

The following table details the options available on the **Initial Run** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows. The settings in this table are only visible if the **Use Cloud-Init/Sysprep** check box is selected, and certain options are only visible when either a Linux-based or Windows-based option has been selected in the **Operating System** list in the **General** tab, as outlined below.

**NOTE**

This table does not include information on whether a power cycle is required because the settings apply to the virtual machine's *initial run*; the virtual machine is not running when you configure these settings.

Table A.3. Virtual Machine: Initial Run Settings

Field Name	Operating System	Description
Use Cloud-Init/Sysprep	Linux, Windows	This check box toggles whether Cloud-Init or Sysprep will be used to initialize the virtual machine.
VM Hostname	Linux, Windows	The host name of the virtual machine.
Domain	Windows	The Active Directory domain to which the virtual machine belongs.
Organization Name	Windows	The name of the organization to which the virtual machine belongs. This option corresponds to the text field for setting the organization name displayed when a machine running Windows is started for the first time.
Active Directory OU	Windows	The organizational unit in the Active Directory domain to which the virtual machine belongs.
Configure Time Zone	Linux, Windows	The time zone for the virtual machine. Select this check box and select a time zone from the Time Zone list.
Admin Password	Windows	<p>The administrative user password for the virtual machine. Click the disclosure arrow to display the settings for this option.</p> <ul style="list-style-type: none"> ● Use already configured password: This check box is automatically selected after you specify an initial administrative user password. You must clear this check box to enable the Admin Password and Verify Admin Password fields and specify a new password. ● Admin Password: The administrative user password for the virtual machine. Enter the password in this text field and the Verify Admin Password text field to verify the password.


Field Name	Operating System	Description
Authentication	Linux	<p>The authentication details for the virtual machine. Click the disclosure arrow to display the settings for this option.</p> <ul style="list-style-type: none"> ● Use already configured password: This check box is automatically selected after you specify an initial root password. You must clear this check box to enable the Password and Verify Password fields and specify a new password. ● Password: The root password for the virtual machine. Enter the password in this text field and the Verify Password text field to verify the password. ● SSH Authorized Keys: SSH keys to be added to the authorized keys file of the virtual machine. You can specify multiple SSH keys by entering each SSH key on a new line. ● Regenerate SSH Keys: Regenerates SSH keys for the virtual machine.
Custom Locale	Windows	<p>Custom locale options for the virtual machine. Locales must be in a format such as en-US. Click the disclosure arrow to display the settings for this option.</p> <ul style="list-style-type: none"> ● Input Locale: The locale for user input. ● UI Language: The language used for user interface elements such as buttons and menus. ● System Locale: The locale for the overall system. ● User Locale: The locale for users.

Field Name	Operating System	Description
Networks	Linux	<p>Network-related settings for the virtual machine. Click the disclosure arrow to display the settings for this option.</p> <ul style="list-style-type: none"> ● DNS Servers: The DNS servers to be used by the virtual machine. ● DNS Search Domains: The DNS search domains to be used by the virtual machine. ● Network: Configures network interfaces for the virtual machine. Select this check box and click + or - to add or remove network interfaces to or from the virtual machine. When you click +, a set of fields becomes visible that can specify whether to use DHCP, and configure an IP address, netmask, and gateway, and specify whether the network interface will start on boot.
Custom Script	Linux	<p>Custom scripts that will be run on the virtual machine when it starts. The scripts entered in this field are custom YAML sections that are added to those produced by the Manager, and allow you to automate tasks such as creating users and files, configuring yum repositories and running commands. For more information on the format of scripts that can be entered in this field, see the Custom Script documentation.</p>
Sysprep	Windows	<p>A custom Sysprep definition. The definition must be in the format of a complete unattended installation answer file. You can copy and paste the default answer files in the <code>/usr/share/ovirt-engine/conf/sysprep/</code> directory on the machine on which the Red Hat Virtualization Manager is installed and alter the fields as required. See Templates for more information.</p>
Ignition 2.3.0	Red Hat Enterprise Linux CoreOS	<p>When Red Hat Enterprise Linux CoreOS is selected as Operating System, this check box toggles whether Ignition will be used to initialize the virtual machine.</p>

A.1.4. Virtual Machine Console Settings Explained

The following table details the options available on the **Console** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.4. Virtual Machine: Console Settings

Field Name	Description	Power cycle required?
Graphical Console Section	A group of settings.	Yes.
Headless Mode	<p>Select this check box if you do not require a graphical console for the virtual machine.</p> <p>When selected, all other fields in the Graphical Console section are disabled. In the VM Portal, the Console icon in the virtual machine's details view is also disabled.</p> <div style="display: flex; align-items: center;">  <div style="margin-left: 10px;"> <p>IMPORTANT</p> <p>See Configuring Headless Machines for more details and prerequisites for using headless mode.</p> </div> </div>	Yes.
Video Type	Defines the graphics device. QXL is the default and supports both graphic protocols. VGA supports only the VNC protocol.	Yes.
Graphics protocol	Defines which display protocol to use. SPICE is the default protocol. VNC is an alternative option. To allow both protocols select SPICE + VNC .	Yes.
VNC Keyboard Layout	Defines the keyboard layout for the virtual machine. This option is only available when using the VNC protocol.	Yes.
USB enabled	<p>Defines SPICE USB redirection. This check box is not selected by default. This option is only available for virtual machines using the SPICE protocol:</p> <ul style="list-style-type: none"> ● Disabled (check box is cleared) - USB controller devices are added according to the devices.usb.controller value in the osinfo-defaults.properties configuration file. The default for all x86 and x86_64 operating systems is piix3-uhci. For ppc64 systems, the default is nec-xhci. ● Enabled (check box is selected) - Enables native KVM/SPICE USB redirection for Linux and Windows virtual machines. Virtual machines do not require any in-guest agents or drivers for native USB. 	Yes.

Field Name	Description	Power cycle required?
Console Disconnect Action	<p>Defines what happens when the console is disconnected. This is only relevant with SPICE and VNC console connections. This setting can be changed while the virtual machine is running but will not take effect until a new console connection is established. Select either:</p> <ul style="list-style-type: none"> ● No action - No action is taken. ● Lock screen - This is the default option. For all Linux machines and for Windows desktops this locks the currently active user session. For Windows servers, this locks the desktop and the currently active user. ● Logout user - For all Linux machines and Windows desktops, this logs out the currently active user session. For Windows servers, the desktop and the currently active user are logged out. ● Shutdown virtual machine - Initiates a graceful virtual machine shutdown. ● Reboot virtual machine - Initiates a graceful virtual machine reboot. 	No.
Monitors	<p>The number of monitors for the virtual machine. This option is only available for virtual desktops using the SPICE display protocol. You can choose 1, 2 or 4. Note that multiple monitors are not supported for Windows systems with WDDMDoD drivers.</p>	Yes.
Smartcard Enabled	<p>Smart cards are an external hardware security feature, most commonly seen in credit cards, but also used by many businesses as authentication tokens. Smart cards can be used to protect Red Hat Virtualization virtual machines. Select or clear the check box to activate and deactivate Smart card authentication for individual virtual machines.</p>	Yes.
Single Sign On method	<p>Enabling Single Sign On allows users to sign into the guest operating system when connecting to a virtual machine from the VM Portal using the Guest Agent.</p> <ul style="list-style-type: none"> ● Disable Single Sign On - Select this option if you do not want the Guest Agent to attempt to sign into the virtual machine. ● Use Guest Agent - Enables Single Sign On to allow the Guest Agent to sign you into the virtual machine. 	If you select Use Guest Agent, no. Otherwise, yes.

Field Name	Description	Power cycle required?
Disable strict user checking	<p>Click the Advanced Parameters arrow and select the check box to use this option. With this option selected, the virtual machine does not need to be rebooted when a different user connects to it.</p> <p>By default, strict checking is enabled so that only one user can connect to the console of a virtual machine. No other user is able to open a console to the same virtual machine until it has been rebooted. The exception is that a SuperUser can connect at any time and replace a existing connection. When a SuperUser has connected, no normal user can connect again until the virtual machine is rebooted.</p> <p>Disable strict checking with caution, because you can expose the previous user's session to the new user.</p>	No.
Soundcard Enabled	A sound card device is not necessary for all virtual machine use cases. If it is for yours, enable a sound card here.	Yes.
Enable SPICE file transfer	Defines whether a user is able to drag and drop files from an external host into the virtual machine's SPICE console. This option is only available for virtual machines using the SPICE protocol. This check box is selected by default.	No.
Enable SPICE clipboard copy and paste	Defines whether a user is able to copy and paste content from an external host into the virtual machine's SPICE console. This option is only available for virtual machines using the SPICE protocol. This check box is selected by default.	No.
Serial Console Section	A group of settings.	
Enable VirtIO serial console	The VirtIO serial console is emulated through VirtIO channels, using SSH and key pairs, and allows you to access a virtual machine's serial console directly from a client machine's command line, instead of opening a console from the Administration Portal or the VM Portal. The serial console requires direct access to the Manager, since the Manager acts as a proxy for the connection, provides information about virtual machine placement, and stores the authentication keys. Select the check box to enable the VirtIO console on the virtual machine. Requires a firewall rule. See Opening a Serial Console to a Virtual Machine .	Yes.

A.1.5. Virtual Machine Host Settings Explained


The following table details the options available on the **Host** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.5. Virtual Machine: Host Settings

Field Name	Sub-element	Description	Power cycle required?
Start Running On		<p>Defines the preferred host on which the virtual machine is to run. Select either:</p> <ul style="list-style-type: none"> ● Any Host in Cluster - The virtual machine can start and run on any available host in the cluster. ● Specific Host(s) - The virtual machine will start running on a particular host in the cluster. However, the Manager or an administrator can migrate the virtual machine to a different host in the cluster depending on the migration and high-availability settings of the virtual machine. Select the specific host or group of hosts from the list of available hosts. 	No. The virtual machine can migrate to that host while running.
CPU options	Pass-Through Host CPU	When selected, allows virtual machines to use the host's CPU flags. When selected, Migration Options is set to Allow manual migration only .	Yes
	Migrate only to hosts with the same TSC frequency	When selected, this virtual machine can only be migrated to a host with the same TSC frequency. This option is only valid for High Performance virtual machines.	Yes

Field Name	Sub-element	Description	Power cycle required?
Migration Options	Migration mode	<p>Defines options to run and migrate the virtual machine. If the options here are not used, the virtual machine will run or migrate according to its cluster's policy.</p> <ul style="list-style-type: none"> ● Allow manual and automatic migration - The virtual machine can be automatically migrated from one host to another in accordance with the status of the environment, or manually by an administrator. ● Allow manual migration only - The virtual machine can only be migrated from one host to another manually by an administrator. ● Do not allow migration - The virtual machine cannot be migrated, either automatically or manually. 	No
	Migration policy	<p>Defines the migration convergence policy. If the check box is left unselected, the host determines the policy.</p> <ul style="list-style-type: none"> ● Cluster default (Minimal downtime) - Overrides in vdsm.conf are still applied. The guest agent hook mechanism is disabled. ● Minimal downtime - Allows the virtual machine to migrate in typical situations. Virtual machines should not experience any significant downtime. The migration will be aborted if virtual machine migration does not converge after a long time (dependent on QEMU iterations, with a maximum of 500 milliseconds). The guest agent hook mechanism is enabled. ● Post-copy migration - When used, post-copy migration pauses the migrating virtual machine vCPUs on the source host, transfers only a minimum of memory pages, activates the virtual machine vCPUs on the destination host, and transfers 	No

Field Name	Sub-element	Description	Power cycle required?
		<p>the remaining memory pages while the virtual machine is running on the destination. The post-copy policy first tries pre-copy to verify whether convergence can occur. The migration switches to post-copy if the virtual machine migration does not converge after a long time.</p> <p>This significantly reduces the downtime of the migrated virtual machine, and also guarantees that the migration finishes regardless of how rapidly the memory pages of the source virtual machine change. It is optimal for migrating virtual machines in heavy continuous use, which would not be possible to migrate with standard pre-copy migration.</p> <p>The disadvantage of this policy is that in the post-copy phase, the virtual machine may slow down significantly as the missing parts of memory are transferred between the hosts.</p>	

Field Name	Sub-element	Description	Power cycle required?
		<div style="text-align: center;">  <p>WARNING</p> </div> <p>If the network connection breaks prior to the completion of the post-copy process, the Manager pauses and then kills the running virtual machine.</p>	
	<p>Enable migration encryption</p>	<p>Allows the virtual machine to be encrypted during migration.</p> <ul style="list-style-type: none"> ● Cluster default ● Encrypt ● Don't encrypt 	<p>No</p>
	<p>Parallel Migrations</p>	<p>Allows you to specify whether and how many parallel migration connections to use.</p> <ul style="list-style-type: none"> ● Cluster default: Parallel migration connections are determined by the cluster default. ● Suspend workload if needed - Disabled: The virtual machine is migrated using a single non-parallel connection. The virtual machine is running a heavy workload. Because of this, the number of parallel connections is automatically determined. This setting might automatically disable all the connections. ● Auto Parallel: The number of parallel connections is automatically determined. ● Custom: Allows you to specify the preferred number of parallel connections, the actual number may be lower. 	


Field Name	Sub-element	Description	Power cycle required?
	Number of VM Migration Connections	This setting is only available when Custom is selected. The preferred number of custom parallel migrations, between 2 and 255.	
Configure NUMA	NUMA Node Count	The number of virtual NUMA nodes available in a host that can be assigned to the virtual machine.	No
	NUMA Pinning	<p>Opens the NUMA Topology window. This window shows the host's total CPUs, memory, and NUMA nodes, and the virtual machine's virtual NUMA nodes. You can manually pin virtual NUMA nodes to host NUMA nodes by clicking and dragging each vNUMA from the box on the right to a NUMA node on the left.</p> <p>You can also set Tune Mode for memory allocation:</p> <p>Strict - Memory allocation will fail if the memory cannot be allocated on the target node.</p> <p>Preferred - Memory is allocated from a single preferred node. If sufficient memory is not available, memory can be allocated from other nodes.</p> <p>Interleave - Memory is allocated across nodes in a round-robin algorithm.</p> <p>If you define NUMA pinning, Migration Options is set to Allow manual migration only.</p>	Yes

A.1.6. Virtual Machine High Availability Settings Explained

The following table details the options available on the **High Availability** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.6. Virtual Machine: High Availability Settings

Field Name	Description	Power cycle required?
------------	-------------	-----------------------

Field Name	Description	Power cycle required?
Highly Available	<p>Select this check box if the virtual machine is to be highly available. For example, in cases of host maintenance, all virtual machines are automatically live migrated to another host. If the host crashes and is in a non-responsive state, only virtual machines with high availability are restarted on another host. If the host is manually shut down by the system administrator, the virtual machine is not automatically live migrated to another host.</p> <p>Note that this option is unavailable for virtual machines defined as Server or Desktop if the Migration Options setting in the Hosts tab is set to Do not allow migration. For a virtual machine to be highly available, it must be possible for the Manager to migrate the virtual machine to other available hosts as necessary.</p> <p>However, for virtual machines defined as High Performance, you can define high availability regardless of the Migration Options setting.</p>	Yes.
Target Storage Domain for VM Lease	<p>Select the storage domain to hold a virtual machine lease, or select No VM Lease to disable the functionality. When a storage domain is selected, it will hold a virtual machine lease on a special volume that allows the virtual machine to be started on another host if the original host loses power or becomes unresponsive.</p> <p>This functionality is only available on storage domain V4 or later.</p> <div data-bbox="488 1406 598 1541" style="display: inline-block; vertical-align: middle;">  </div> <p>NOTE</p> <p>If you define a lease, the only Resume Behavior available is KILL.</p>	Yes.

Field Name	Description	Power cycle required?
Resume Behavior	<p>Defines the desired behavior of a virtual machine that is paused due to storage I/O errors, once a connection with the storage is reestablished. You can define the desired resume behavior even if the virtual machine is not highly available.</p> <p>The following options are available:</p> <ul style="list-style-type: none"> ● AUTO_RESUME - The virtual machine is automatically resumed, without requiring user intervention. This is recommended for virtual machines that are not highly available and that do not require user intervention after being in the paused state. ● LEAVE_PAUSED - The virtual machine remains in pause mode until it is manually resumed or restarted. ● KILL - The virtual machine is automatically resumed if the I/O error is remedied within 80 seconds. However, if more than 80 seconds pass, the virtual machine is ungracefully shut down. This is recommended for highly available virtual machines, to allow the Manager to restart them on another host that is not experiencing the storage I/O error. KILL is the only option available when using virtual machine leases. 	No.
Priority for Run/Migration queue	Sets the priority level for the virtual machine to be migrated or restarted on another host.	No.

Field Name	Description	Power cycle required?
Watchdog	<p>Allows users to attach a watchdog card to a virtual machine. A watchdog is a timer that is used to automatically detect and recover from failures. Once set, a watchdog timer continually counts down to zero while the system is in operation, and is periodically restarted by the system to prevent it from reaching zero. If the timer reaches zero, it signifies that the system has been unable to reset the timer and is therefore experiencing a failure. Corrective actions are then taken to address the failure. This functionality is especially useful for servers that demand high availability.</p> <p>Watchdog Model: The model of watchdog card to assign to the virtual machine. At current, the only supported model is i6300esb.</p> <p>Watchdog Action: The action to take if the watchdog timer reaches zero. The following actions are available:</p> <ul style="list-style-type: none"> ● none - No action is taken. However, the watchdog event is recorded in the audit log. ● reset - The virtual machine is reset and the Manager is notified of the reset action. ● poweroff - The virtual machine is immediately shut down. ● dump - A dump is performed and the virtual machine is paused. The guest's memory is dumped by libvirt, therefore, neither 'kdump' nor 'pvpanic' is required. The dump file is created in the directory that is configured by <code>auto_dump_path</code> in the <code>/etc/libvirt/qemu.conf</code> file on the host. ● pause - The virtual machine is paused, and can be resumed by users. 	Yes.

A.1.7. Virtual Machine Resource Allocation Settings Explained

The following table details the options available on the **Resource Allocation** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.7. Virtual Machine: Resource Allocation Settings

Field Name	Sub-element	Description	Power cycle required?
------------	-------------	-------------	-----------------------

Field Name	Sub-element	Description	Power cycle required?
CPU Allocation	CPU Profile	The CPU profile assigned to the virtual machine. CPU profiles define the maximum amount of processing capability a virtual machine can access on the host on which it runs, expressed as a percent of the total processing capability available to that host. CPU profiles are defined for a cluster, based on quality of service entries created for data centers.	No.
	CPU Shares	Allows users to set the level of CPU resources a virtual machine can demand relative to other virtual machines. <ul style="list-style-type: none">● Low - 512● Medium - 1024● High - 2048● Custom - A custom level of CPU shares defined by the user.	No.

Field Name	Sub-element	Description	Power cycle required?
	CPU Pinning Policy	<ul style="list-style-type: none">● None - Runs without any CPU pinning.● Manual - Runs a manually specified virtual CPU on a specific physical CPU and a specific host. Available only when the virtual machine is pinned to a Host.● Resize and Pin NUMA - Resizes the virtual CPU and NUMA topology of the virtual machine according to the Host, and pins them to the Host resources.● Dedicated - Exclusively pins virtual CPUs to host physical CPUs. Available for cluster compatibility level 4.7 or later. If the virtual machine has NUMA enabled, all nodes must be unpinned.● Isolate Threads - Exclusively pins virtual CPUs to host physical CPUs. Each virtual CPU gets a physical core. Available for cluster compatibility level 4.7 or later. If the virtual machine has NUMA enabled, all nodes must be unpinned.	No.

Field Name	Sub-element	Description	Power cycle required?
	CPU Pinning topology	<p>Enables the virtual machine's virtual CPU (vCPU) to run on a specific physical CPU (pCPU) in a specific host. The syntax of CPU pinning is v#p[_v#p], for example:</p> <ul style="list-style-type: none"> ● 0#0 - Pins vCPU 0 to pCPU 0. ● 0#0_1#3 - Pins vCPU 0 to pCPU 0, and pins vCPU 1 to pCPU 3. ● 1#1-4,^2 - Pins vCPU 1 to one of the pCPUs in the range of 1 to 4, excluding pCPU 2. <p>The CPU Pinning Topology is populated automatically when Resize and Pin NUMA pinning is selected in CPU Pinning Policy.</p> <p>In order to pin a virtual machine to a host, you must also select the following on the Host tab:</p> <ul style="list-style-type: none"> ● Start Running On: Specific ● Pass-Through Host CPU <p>If CPU pinning is set and you change Start Running On: Specific a CPU pinning topology will be lost window appears when you click OK.</p> <p>When defined, Migration Options in the Hosts tab is set to Allow manual migration only.</p>	Yes.
Memory Allocation	Physical Memory Guaranteed	<p>The amount of physical memory guaranteed for this virtual machine. Should be any number between 0 and the defined memory for this virtual machine.</p>	If lowered, yes. Otherwise, no.

Field Name	Sub-element	Description	Power cycle required?
	Memory Balloon Device Enabled	Enables the memory balloon device for this virtual machine. Enable this setting to allow memory overcommitment in a cluster. Enable this setting for applications that allocate large amounts of memory suddenly but set the guaranteed memory to the same value as the defined memory. Use ballooning for applications and loads that slowly consume memory, occasionally release memory, or stay dormant for long periods of time, such as virtual desktops. See Optimization Settings Explained in the <i>Administration Guide</i> for more information.	Yes.
Trusted Platform Module	TPM Device Enabled	Enables the addition of an emulated Trusted Platform Module (TPM) device. Select this check box to add an emulated Trusted Platform Module device to a virtual machine. TPM devices can only be used on x86_64 machines with UEFI firmware and PowerPC machines with pSeries firmware installed. See Adding Trusted Platform Module devices for more information.	Yes.
IO Threads	IO Threads Enabled	Enables IO threads. Select this check box to improve the speed of disks that have a VirtIO interface by pinning them to a thread separate from the virtual machine's other functions. Improved disk performance increases a virtual machine's overall performance. Disks with VirtIO interfaces are pinned to an IO thread using a round-robin algorithm.	Yes.

Field Name	Sub-element	Description	Power cycle required?
Queues	Multi Queues Enabled	<p>Enables multiple queues. This check box is selected by default. It creates up to four queues per vNIC, depending on how many vCPUs are available.</p> <p>It is possible to define a different number of queues per vNIC by creating a custom property as follows:</p> <p>engine-config -s "CustomDeviceProperties={type=interface;prop={<i>other-nic-properties</i>;queues=[1-9][0-9]*}"</p> <p>where <i>other-nic-properties</i> is a semicolon-separated list of pre-existing NIC custom properties.</p>	Yes.
	VirtIO-SCSI Enabled	Allows users to enable or disable the use of VirtIO-SCSI on the virtual machines.	Not applicable.
	VirtIO-SCSI Multi Queues Enabled	The VirtIO-SCSI Multi Queues Enabled option is only available when VirtIO-SCSI Enabled is selected. Select this check box to enable multiple queues in the VirtIO-SCSI driver. This setting can improve I/O throughput when multiple threads within the virtual machine access the virtual disks. It creates up to four queues per VirtIO-SCSI controller, depending on how many disks are connected to the controller and how many vCPUs are available.	Not applicable.
Storage Allocation		The Storage Allocation option is only available when the virtual machine is created from a template.	Not applicable.
	Thin	Provides optimized usage of storage capacity. Disk space is allocated only as it is required. When selected, the format of the disks will be marked as QCOW2 and you will not be able to change it.	Not applicable.

Field Name	Sub-element	Description	Power cycle required?
	Clone	Optimized for the speed of guest read and write operations. All disk space requested in the template is allocated at the time of the clone operation. Possible disk formats are QCOW2 or Raw .	Not applicable.
Disk Allocation		The Disk Allocation option is only available when you are creating a virtual machine from a template.	Not applicable.
	Alias	An alias for the virtual disk. By default, the alias is set to the same value as that of the template.	Not applicable.
	Virtual Size	The total amount of disk space that the virtual machine based on the template can use. This value cannot be edited, and is provided for reference only.	Not applicable.
	Format	The format of the virtual disk. The available options are QCOW2 and Raw . When Storage Allocation is Thin , the disk format is QCOW2 . When Storage Allocation is Clone , select QCOW2 or Raw .	Not applicable.
	Target	The storage domain on which the virtual disk is stored. By default, the storage domain is set to the same value as that of the template.	Not applicable.
	Disk Profile	The disk profile to assign to the virtual disk. Disk profiles are created based on storage profiles defined in the data centers. For more information, see Creating a Disk Profile .	Not applicable.

A.1.8. Virtual Machine Boot Options Settings Explained

The following table details the options available on the **Boot Options** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows

Table A.8. Virtual Machine: Boot Options Settings

Field Name	Description	Power cycle required?
------------	-------------	-----------------------

Field Name	Description	Power cycle required?
First Device	<p>After installing a new virtual machine, the new virtual machine must go into Boot mode before powering up. Select the first device that the virtual machine must try to boot:</p> <ul style="list-style-type: none"> ● Hard Disk ● CD-ROM ● Network (PXE) 	Yes.
Second Device	Select the second device for the virtual machine to use to boot if the first device is not available. The first device selected in the previous option does not appear in the options.	Yes.
Attach CD	If you have selected CD-ROM as a boot device, select this check box and select a CD-ROM image from the drop-down menu. The images must be available in the ISO domain.	Yes.
Enable menu to select boot device	Enables a menu to select the boot device. After the virtual machine starts and connects to the console, but before the virtual machine starts booting, a menu displays that allows you to select the boot device. This option should be enabled before the initial boot to allow you to select the required installation media.	Yes.

A.1.9. Virtual Machine Random Generator Settings Explained

The following table details the options available on the **Random Generator** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.9. Virtual Machine: Random Generator Settings

Field Name	Description	Power cycle required?
Random Generator enabled	Selecting this check box enables a paravirtualized Random Number Generator PCI device (virtio-rng). This device allows entropy to be passed from the host to the virtual machine in order to generate a more sophisticated random number. Note that this check box can only be selected if the RNG device exists on the host and is enabled in the host's cluster.	Yes.
Period duration (ms)	Specifies the duration of the RNG's "full cycle" or "full period" in milliseconds. If omitted, the libvirt default of 1000 milliseconds (1 second) is used. If this field is filled, Bytes per period must be filled also.	Yes.

Field Name	Description	Power cycle required?
Bytes per period	Specifies how many bytes are permitted to be consumed per period.	Yes.
Device source:	<p>The source of the random number generator. This is automatically selected depending on the source supported by the host's cluster.</p> <ul style="list-style-type: none"> ● /dev/urandom source - The Linux-provided random number generator. ● /dev/hwrng source - An external hardware generator. 	Yes.

A.1.10. Virtual Machine Custom Properties Settings Explained

The following table details the options available on the **Custom Properties** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.

Table A.10. Virtual Machine Custom Properties Settings

Field Name	Description	Recommendations and Limitations	Power cycle required?
sndbuf	Enter the size of the buffer for sending the virtual machine's outgoing data over the socket. Default value is 0.	-	Yes
hugepages	Enter the huge page size in KB.	<ul style="list-style-type: none"> ● Set the huge page size to the largest size supported by the pinned host. ● The recommended size for x86_64 is 1 GB. ● The virtual machine's huge page size must be the same size as the pinned host's huge page size. ● The virtual machine's memory size must fit into the selected size of the pinned host's free huge pages. The NUMA node size must be a multiple of the huge page's selected size. 	Yes

Field Name	Description	Recommendations and Limitations	Power cycle required?
vhost	<p>Disables vhost-net, which is the kernel-based virtio network driver on virtual network interface cards attached to the virtual machine. To disable vhost, the format for this property is <i>LogicalNetworkName</i>: false.</p> <p>This will explicitly start the virtual machine without the vhost-net setting on the virtual NIC attached to <i>LogicalNetworkName</i>.</p>	vhost-net provides better performance than virtio-net, and if it is present, it is enabled on all virtual machine NICs by default. Disabling this property makes it easier to isolate and diagnose performance issues, or to debug vhost-net errors; for example, if migration fails for virtual machines on which vhost does not exist.	Yes
sap_agent	Enables SAP monitoring on the virtual machine. Set to true or false .	-	Yes
viodiskcache	Caching mode for the virtio disk. writethrough writes data to the cache and the disk in parallel, writeback does not copy modifications from the cache to the disk, and none disables caching.	In order to ensure data integrity in the event of a fault in storage, in the network, or in a host during migration, do not migrate virtual machines with viodiskcache enabled, unless virtual machine clustering or application-level clustering is also enabled.	Yes

Field Name	Description	Recommendations and Limitations	Power cycle required?
<code>scsi_hostdev</code>	<p>Optionally, if you add a SCSI host device to a virtual machine, you can specify the optimal SCSI host device driver. For details, see Adding Host Devices to a Virtual Machine.</p> <ul style="list-style-type: none"> ● scsi_generic: (Default) Enables the guest operating system to access OS-supported SCSI host devices attached to the host. Use this driver for SCSI media changers that require raw access, such as tape or CD changers. ● scsi_block: Similar to scsi_generic but better speed and reliability. Use for SCSI disk devices. If trim or discard for the underlying device is desired, and it's a hard disk, use this driver. ● scsi_hd: Provides performance with lowered overhead. Supports large numbers of devices. Uses the standard SCSI device naming scheme. Can be used with aio-native. Use this driver for high-performance SSDs. ● virtio_blk_pci: Provides the highest performance without the SCSI overhead. Supports identifying devices by their serial numbers. 	If you are not sure, try <code>scsi_hd</code> .	Yes



WARNING

Increasing the value of the `sndbuf` custom property results in increased occurrences of communication failure between hosts and unresponsive virtual machines.

A.1.11. Virtual Machine Icon Settings Explained

You can add custom icons to virtual machines and templates. Custom icons can help to differentiate virtual machines in the VM Portal. The following table details the options available on the **Icon** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows.



NOTE

This table does not include information on whether a power cycle is required because these settings apply to the virtual machine's *appearance in the Administration portal*, not to its configuration.

Table A.11. Virtual Machine: Icon Settings

Button Name	Description
Upload	Click this button to select a custom image to use as the virtual machine's icon. The following limitations apply: <ul style="list-style-type: none"> • Supported formats: jpg, png, gif • Maximum size: 24 KB • Maximum dimensions: 150px width, 120px height
Power cycle required?	Use default

A.1.12. Virtual Machine Foreman/Satellite Settings Explained

The following table details the options available on the **Foreman/Satellite** tab of the **New Virtual Machine** and **Edit Virtual Machine** windows

Table A.12. Virtual Machine:Foreman/Satellite Settings

Field Name	Description	Power cycle required?
Provider	If the virtual machine is running Red Hat Enterprise Linux and the system is configured to work with a Satellite server, select the name of the Satellite from the list. This enables you to use Satellite's content management feature to display the relevant Errata for this virtual machine. See Configuring Satellite Errata for more details.	Yes.


A.2. EXPLANATION OF SETTINGS IN THE RUN ONCE WINDOW

The **Run Once** window defines one-off boot options for a virtual machine. For persistent boot options, use the **Boot Options** tab in the **New Virtual Machine** window. The **Run Once** window contains multiple sections that can be configured.

The standalone **Rollback this configuration during reboots** check box specifies whether reboots (initiated by the Manager, or from within the guest) will be warm (soft) or cold (hard). Select this check

box to configure a cold reboot that restarts the virtual machine with regular (non-**Run Once**) configuration. Clear this check box to configure a warm reboot that retains the virtual machine's **Run Once** configuration.

The **Boot Options** section defines the virtual machine's boot sequence, running options, and source images for installing the operating system and required drivers.



NOTE

The following tables do not include information on whether a power cycle is required because these *one-off boot options* apply only when you reboot the virtual machine.

Table A.13. Boot Options Section

Field Name	Description
Attach CD	Attaches an ISO image to the virtual machine. Use this option to install the virtual machine's operating system and applications. The CD image must reside in the ISO domain.
Attach Windows guest tools CD	Attaches a secondary virtual CD-ROM to the virtual machine with the virtio-win ISO image. Use this option to install Windows drivers. For information on installing the image, see Uploading the VirtIO Image Files to a Storage Domain in the <i>Administration Guide</i> .
Enable menu to select boot device	Enables a menu to select the boot device. After the virtual machine starts and connects to the console, but before the virtual machine starts booting, a menu displays that allows you to select the boot device. This option should be enabled before the initial boot to allow you to select the required installation media.
Start in Pause Mode	Starts and then pauses the virtual machine to enable connection to the console. Suitable for virtual machines in remote locations.
Predefined Boot Sequence	Determines the order in which the boot devices are used to boot the virtual machine. Select Hard Disk , CD-ROM , or Network (PXE) , and use Up and Down to move the option up or down in the list.
Run Stateless	Deletes all data and configuration changes to the virtual machine upon shutdown. This option is only available if a virtual disk is attached to the virtual machine.

The **Linux Boot Options** section contains fields to boot a Linux kernel directly instead of through the BIOS bootloader.

Table A.14. Linux Boot Options Section

Field Name	Description
------------	-------------

Field Name	Description
kernel path	A fully qualified path to a kernel image to boot the virtual machine. The kernel image must be stored on either the ISO domain (path name in the format of iso://path-to-image) or on the host's local storage domain (path name in the format of /data/images).
initrd path	A fully qualified path to a ramdisk image to be used with the previously specified kernel. The ramdisk image must be stored on the ISO domain (path name in the format of iso://path-to-image) or on the host's local storage domain (path name in the format of /data/images).
kernel parameters	Kernel command line parameter strings to be used with the defined kernel on boot.

The **Initial Run** section is used to specify whether to use Cloud-Init or Sysprep to initialize the virtual machine. For Linux-based virtual machines, you must select the **Use Cloud-Init** check box in the **Initial Run** tab to view the available options. For Windows-based virtual machines, you must attach the **[sysprep]** floppy by selecting the **Attach Floppy** check box in the **Boot Options** tab and selecting the floppy from the list.

The options that are available in the **Initial Run** section differ depending on the operating system that the virtual machine is based on.

Table A.15. Initial Run Section (Linux-based Virtual Machines)

Field Name	Description
VM Hostname	The host name of the virtual machine.
Configure Time Zone	The time zone for the virtual machine. Select this check box and select a time zone from the Time Zone list.
Authentication	The authentication details for the virtual machine. Click the disclosure arrow to display the settings for this option.
Authentication → User Name	Creates a new user account on the virtual machine. If this field is not filled in, the default user is root .
Authentication → Use already configured password	This check box is automatically selected after you specify an initial root password. You must clear this check box to enable the Password and Verify Password fields and specify a new password.
Authentication → Password	The root password for the virtual machine. Enter the password in this text field and the Verify Password text field to verify the password.
Authentication → SSH Authorized Keys	SSH keys to be added to the authorized keys file of the virtual machine.

Field Name	Description
Authentication → Regenerate SSH Keys	Regenerates SSH keys for the virtual machine.
Networks	Network-related settings for the virtual machine. Click the disclosure arrow to display the settings for this option.
Networks → DNS Servers	The DNS servers to be used by the virtual machine.
Networks → DNS Search Domains	The DNS search domains to be used by the virtual machine.
Networks → Network	Configures network interfaces for the virtual machine. Select this check box and click + or - to add or remove network interfaces to or from the virtual machine. When you click +, a set of fields becomes visible that can specify whether to use DHCP, and configure an IP address, netmask, and gateway, and specify whether the network interface will start on boot.
Custom Script	Custom scripts that will be run on the virtual machine when it starts. The scripts entered in this field are custom YAML sections that are added to those produced by the Manager, and allow you to automate tasks such as creating users and files, configuring yum repositories and running commands. For more information on the format of scripts that can be entered in this field, see the Custom Script documentation.

Table A.16. Initial Run Section (Windows-based Virtual Machines)

Field Name	Description
VM Hostname	The host name of the virtual machine.
Domain	The Active Directory domain to which the virtual machine belongs.
Organization Name	The name of the organization to which the virtual machine belongs. This option corresponds to the text field for setting the organization name displayed when a machine running Windows is started for the first time.
Active Directory OU	The organizational unit in the Active Directory domain to which the virtual machine belongs. The distinguished name must be provided. For example CN=Users,DC=lab,DC=local
Configure Time Zone	The time zone for the virtual machine. Select this check box and select a time zone from the Time Zone list.
Admin Password	The administrative user password for the virtual machine. Click the disclosure arrow to display the settings for this option.

Field Name	Description
Admin Password → Use already configured password	This check box is automatically selected after you specify an initial administrative user password. You must clear this check box to enable the Admin Password and Verify Admin Password fields and specify a new password.
Admin Password → Admin Password	The administrative user password for the virtual machine. Enter the password in this text field and the Verify Admin Password text field to verify the password.
Custom Locale	Locales must be in a format such as en-US . Click the disclosure arrow to display the settings for this option.
Custom Locale → Input Locale	The locale for user input.
Custom Locale → UI Language	The language used for user interface elements such as buttons and menus.
Custom Locale → System Locale	The locale for the overall system.
Custom Locale → User Locale	The locale for users.
Sysprep	A custom Sysprep definition. The definition must be in the format of a complete unattended installation answer file. You can copy and paste the default answer files in the <code>/usr/share/ovirt-engine/conf/sysprep/</code> directory on the machine on which the Red Hat Virtualization Manager is installed and alter the fields as required. The definition will overwrite any values entered in the Initial Run fields. See Templates for more information.
Domain	The Active Directory domain to which the virtual machine belongs. If left blank, the value of the previous Domain field is used.
Alternate Credentials	Selecting this check box allows you to set a User Name and Password as alternative credentials.

The **System** section enables you to define the supported machine type or CPU type.

Table A.17. System Section

Field Name	Description
Custom Emulated Machine	This option allows you to specify the machine type. If changed, the virtual machine will only run on hosts that support this machine type. Defaults to the cluster's default machine type.

Field Name	Description
Custom CPU Type	This option allows you to specify a CPU type. If changed, the virtual machine will only run on hosts that support this CPU type. Defaults to the cluster's default CPU type.

The **Host** section is used to define the virtual machine's host.

Table A.18. Host Section

Field Name	Description
Any host in cluster	Allocates the virtual machine to any available host.
Specific Host(s)	Specifies a user-defined host for the virtual machine.

The **Console** section defines the protocol to connect to virtual machines.

Table A.19. Console Section

Field Name	Description
Headless Mode	Select this option if you do not require a graphical console when running the machine for the first time. See Configuring Headless Machines for more information.
VNC	Requires a VNC client to connect to a virtual machine using VNC. Optionally, specify VNC Keyboard Layout from the drop-down list.
SPICE	Recommended protocol for Linux and Windows virtual machines. Using SPICE protocol with QXLDDOD drivers is supported for Windows 10 and Windows Server 2016 and later virtual machines.
Enable SPICE file transfer	Determines whether you can drag and drop files from an external host into the virtual machine's SPICE console. This option is only available for virtual machines using the SPICE protocol. This check box is selected by default.
Enable SPICE clipboard copy and paste	Defines whether you can copy and paste content from an external host into the virtual machine's SPICE console. This option is only available for virtual machines using the SPICE protocol. This check box is selected by default.

The **Custom Properties** section contains additional VDSM options for running virtual machines. See [New VMs Custom Properties](#) for details.

A.3. EXPLANATION OF SETTINGS IN THE NEW NETWORK INTERFACE AND EDIT NETWORK INTERFACE WINDOWS

These settings apply when you are adding or editing a virtual machine network interface. If you have more than one network interface attached to a virtual machine, you can put the virtual machine on more than one logical network.

Table A.20. Network Interface Settings

Field Name	Description	Power cycle required?
Name	The name of the network interface. This text field has a 21-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.	No.
Profile	The vNIC profile and logical network that the network interface is placed on. By default, all network interfaces are put on the ovirtmgmt management network.	No.
Type	<p>The virtual interface the network interface presents to virtual machines.</p> <ul style="list-style-type: none"> ● rtl8139 and e1000 device drivers are included in most operating systems. ● VirtIO is faster but requires VirtIO drivers. Red Hat Enterprise Linux 5 and later include VirtIO drivers. Windows does not include VirtIO drivers, but they can be installed from the guest tools ISO or virtual floppy disk. ● PCI Passthrough enables the vNIC to be directly connected to a virtual function (VF) of an SR-IOV-enabled NIC. The vNIC will then bypass the software network virtualization and connect directly to the VF for direct device assignment. The selected vNIC profile must also have Passthrough enabled. 	Yes.
Custom MAC address	Choose this option to set a custom MAC address. The Red Hat Virtualization Manager automatically generates a MAC address that is unique to the environment to identify the network interface. Having two devices with the same MAC address online in the same network causes networking conflicts.	Yes.

Field Name	Description	Power cycle required?
Link State	<p>Whether or not the network interface is connected to the logical network.</p> <ul style="list-style-type: none"> ● Up: The network interface is located on its slot. <ul style="list-style-type: none"> ○ When the Card Status is Plugged, it means the network interface is connected to a network cable, and is active. ○ When the Card Status is Unplugged, the network interface will automatically be connected to the network and become active once plugged. ● Down: The network interface is located on its slot, but it is not connected to any network. Virtual machines will not be able to run in this state. 	No.
Card Status	<p>Whether or not the network interface is defined on the virtual machine.</p> <ul style="list-style-type: none"> ● Plugged: The network interface has been defined on the virtual machine. <ul style="list-style-type: none"> ○ If its Link State is Up, it means the network interface is connected to a network cable, and is active. ○ If its Link State is Down, the network interface is not connected to a network cable. ● Unplugged: The network interface is only defined on the Manager, and is not associated with a virtual machine. <ul style="list-style-type: none"> ○ If its Link State is Up, when the network interface is plugged it will automatically be connected to a network and become active. ○ If its Link State is Down, the network interface is not connected to any network until it is defined on a virtual machine. 	No.

A.4. EXPLANATION OF SETTINGS IN THE NEW VIRTUAL DISK AND EDIT VIRTUAL DISK WINDOWS

**NOTE**

The following tables do not include information on whether a power cycle is required because that information is not applicable to these scenarios.

Table A.21. New Virtual Disk and Edit Virtual Disk settings

Field Name	Description
Size(GB)	The size of the new virtual disk in GB.
Alias	The name of the virtual disk, limited to 40 characters.
Description	A description of the virtual disk. This field is recommended but not mandatory.
Interface	<p>The virtual interface the disk presents to virtual machines. VirtIO is faster, but requires drivers. Red Hat Enterprise Linux 5 and later include these drivers. Windows does not include these drivers, but you can install them from the virtio-win ISO image. IDE and SATA devices do not require special drivers.</p> <p>The interface type can be updated after stopping all virtual machines that the disk is attached to.</p>
Data Center	The data center in which the virtual disk will be available.
Storage Domain	The storage domain in which the virtual disk will be stored. The drop-down list shows all storage domains available in the given data center, and also shows the total space and currently available space in the storage domain.
Allocation Policy	<p>The provisioning policy for the new virtual disk.</p> <ul style="list-style-type: none"> ● Preallocated allocates the entire size of the disk on the storage domain at the time the virtual disk is created. The virtual size and the actual size of a preallocated disk are the same. Preallocated virtual disks take more time to create than thin provisioned virtual disks, but have better read and write performance. Preallocated virtual disks are recommended for servers and other I/O intensive virtual machines. If a virtual machine is able to write more than 1 GB every four seconds, use preallocated disks where possible. ● Thin Provision allocates 1 GB at the time the virtual disk is created and sets a maximum limit on the size to which the disk can grow. The virtual size of the disk is the maximum limit; the actual size of the disk is the space that has been allocated so far. Thin provisioned disks are faster to create than preallocated disks and allow for storage over-commitment. Thin provisioned virtual disks are recommended for desktops.

Field Name	Description
Disk Profile	The disk profile assigned to the virtual disk. Disk profiles define the maximum amount of throughput and the maximum level of input and output operations for a virtual disk in a storage domain. Disk profiles are defined on the storage domain level based on storage quality of service entries created for data centers.
Activate Disk(s)	Activate the virtual disk immediately after creation. This option is not available when creating a floating disk.
Wipe After Delete	Allows you to enable enhanced security for deletion of sensitive material when the virtual disk is deleted.
Bootable	Enables the bootable flag on the virtual disk.
Shareable	Attaches the virtual disk to more than one virtual machine at a time.
Read Only	Allows you to set the disk as read-only. The same disk can be attached as read-only to one virtual machine, and as rewritable to another. This option is not available when creating a floating disk.
Enable Discard	Allows you to shrink a thin provisioned disk while the virtual machine is up. For block storage, the underlying storage device must support discard calls, and the option cannot be used with Wipe After Delete unless the underlying storage supports the <code>discard_zeroes_data</code> property. For file storage, the underlying file system and the block device must support discard calls. If all requirements are met, SCSI UNMAP commands issued from guest virtual machines is passed on by QEMU to the underlying storage to free up the unused space.

The **Direct LUN** settings can be displayed in either **Targets > LUNs** or **LUNs > Targets**. **Targets > LUNs** sorts available LUNs according to the host on which they are discovered, whereas **LUNs > Targets** displays a single list of LUNs.

Table A.22. New Virtual Disk and Edit Virtual Disk settingsDirect LUN

Field Name	Description
Alias	The name of the virtual disk, limited to 40 characters.
Description	<p>A description of the virtual disk. This field is recommended but not mandatory. By default the last 4 characters of the LUN ID is inserted into the field.</p> <p>The default behavior can be configured by setting the PopulateDirectLUNDiskDescriptionWithLUNId configuration key to the appropriate value using the engine-config command. The configuration key can be set to -1 for the full LUN ID to be used, or 0 for this feature to be ignored. A positive integer populates the description with the corresponding number of characters of the LUN ID.</p>

Field Name	Description
Interface	<p>The virtual interface the disk presents to virtual machines. VirtIO is faster, but requires drivers. Red Hat Enterprise Linux 5 and later include these drivers. Windows does not include these drivers, but you can install them from the virtio-win ISO image. IDE and SATA devices do not require special drivers.</p> <p>The interface type can be updated after stopping all virtual machines that the disk is attached to.</p>
Data Center	The data center in which the virtual disk will be available.
Host	The host on which the LUN will be mounted. You can select any host in the data center.
Storage Type	The type of external LUN to add. You can select from either iSCSI or Fibre Channel .
Discover Targets	<p>This section can be expanded when you are using iSCSI external LUNs and Targets > LUNs is selected.</p> <p>Address - The host name or IP address of the target server.</p> <p>Port - The port by which to attempt a connection to the target server. The default port is 3260.</p> <p>User Authentication - The iSCSI server requires User Authentication. The User Authentication field is visible when you are using iSCSI external LUNs.</p> <p>CHAP user name - The user name of a user with permission to log in to LUNs. This field is accessible when the User Authentication check box is selected.</p> <p>CHAP password - The password of a user with permission to log in to LUNs. This field is accessible when the User Authentication check box is selected.</p>
Activate Disk(s)	Activate the virtual disk immediately after creation. This option is not available when creating a floating disk.
Bootable	Allows you to enable the bootable flag on the virtual disk.
Shareable	Allows you to attach the virtual disk to more than one virtual machine at a time.
Read Only	Allows you to set the disk as read-only. The same disk can be attached as read-only to one virtual machine, and as rewritable to another. This option is not available when creating a floating disk.

Field Name	Description
Enable Discard	Allows you to shrink a thin provisioned disk while the virtual machine is up. With this option enabled, SCSI UNMAP commands issued from guest virtual machines is passed on by QEMU to the underlying storage to free up the unused space.
Enable SCSI Pass-Through	Available when the Interface is set to VirtIO-SCSI . Selecting this check box enables passthrough of a physical SCSI device to the virtual disk. A VirtIO-SCSI interface with SCSI passthrough enabled automatically includes SCSI discard support. Read Only is not supported when this check box is selected. When this check box is not selected, the virtual disk uses an emulated SCSI device. Read Only is supported on emulated VirtIO-SCSI disks.
Allow Privileged SCSI I/O	Available when the Enable SCSI Pass-Through check box is selected. Selecting this check box enables unfiltered SCSI Generic I/O (SG_IO) access, allowing privileged SG_IO commands on the disk. This is required for persistent reservations.
Using SCSI Reservation	Available when the Enable SCSI Pass-Through and Allow Privileged SCSI I/O check boxes are selected. Selecting this check box disables migration for any virtual machine using this disk, to prevent virtual machines that are using SCSI reservation from losing access to the disk.

Fill in the fields in the **Discover Targets** section and click **Discover** to discover the target server. You can then click the **Login All** button to list the available LUNs on the target server and, using the radio buttons next to each LUN, select the LUN to add.

Using LUNs directly as virtual machine hard disk images removes a layer of abstraction between your virtual machines and their data.

The following considerations must be made when using a direct LUN as a virtual machine hard disk image:

- Live storage migration of direct LUN hard disk images is not supported.
- Direct LUN disks are not included in virtual machine exports.
- Direct LUN disks are not included in virtual machine snapshots.



IMPORTANT

Mounting a journaled file system requires read-write access. Using the **Read Only** option is not appropriate for virtual disks that contain such file systems (e.g. **EXT3**, **EXT4**, or **XFS**).

A.5. EXPLANATION OF SETTINGS IN THE NEW TEMPLATE WINDOW

The following table details the settings for the **New Template** window.

**NOTE**

The following tables do not include information on whether a power cycle is required because that information is not applicable to this scenario.

1. New Template Settings

Field
Description/Action
Name
The name of the template. This is the name by which the template is listed in the Templates tab in the Administration Portal and is accessed via the REST API. This text field has a 40-character limit and must be a unique name within the data center with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores. The name can be reused in different data centers in the environment.
Description
A description of the template. This field is recommended but not mandatory.
Comment
A field for adding plain text, human-readable comments regarding the template.
Cluster
The cluster with which the template is associated. This is the same as the original virtual machines by default. You can select any cluster in the data center.
CPU Profile
The CPU profile assigned to the template. CPU profiles define the maximum amount of processing capability a virtual machine can access on the host on which it runs, expressed as a percent of the total processing capability available to that host. CPU profiles are defined for a cluster based on quality of service entries created for data centers.
Create as a Template Sub-Version
Specifies whether the template is created as a new version of an existing template. Select this check box to access the settings for configuring this option. <ul style="list-style-type: none"> ● Root Template: The template under which the sub-template is added. ● Sub-Version Name: The name of the template. This is the name by which the template is accessed when creating a new virtual machine based on the template. If the virtual machine is stateless, the list of sub-versions will contain a latest option rather than the name of the latest sub-version. This option automatically applies the latest template sub-version to the virtual machine upon reboot. Sub-versions are particularly useful when working with pools of stateless virtual machines.

Disks Allocation

Alias - An alias for the virtual disk used by the template. By default, the alias is set to the same value as that of the source virtual machine.

Virtual Size - The total amount of disk space that a virtual machine based on the template can use. This value cannot be edited, and is provided for reference only. This value corresponds with the size, in GB, that was specified when the disk was created or edited.

Format - The format of the virtual disk used by the template. The available options are QCOW2 and Raw. By default, the format is set to Raw.

Target - The storage domain on which the virtual disk used by the template is stored. By default, the storage domain is set to the same value as that of the source virtual machine. You can select any storage domain in the cluster.

Disk Profile - The disk profile to assign to the virtual disk used by the template. Disk profiles are created based on storage profiles defined in the data centers. For more information, see [Creating a Disk Profile](#).

Allow all users to access this Template

Specifies whether a template is public or private. A public template can be accessed by all users, whereas a private template can only be accessed by users with the **TemplateAdmin** or **SuperUser** roles.

Copy VM permissions

Copies explicit permissions that have been set on the source virtual machine to the template.

Seal Template (Linux only)

Specifies whether a template is sealed. 'Sealing' is an operation that erases all machine-specific configurations from a filesystem, including SSH keys, UDEV rules, MAC addresses, system ID, and hostname. This setting prevents a virtual machine based on this template from inheriting the configuration of the source virtual machine.

APPENDIX B. VIRT-SYSPREP OPERATIONS

The **virt-sysprep** command removes system-specific details.

Only operations marked with * are performed during the template sealing process.

```
# virt-sysprep --list-operations
abrt-data * Remove the crash data generated by ABRT
bash-history * Remove the bash history in the guest
blkid-tab * Remove blkid tab in the guest
ca-certificates Remove CA certificates in the guest
crash-data * Remove the crash data generated by kexec-tools
cron-spool * Remove user at-jobs and cron-jobs
customize * Customize the guest
dhcp-client-state * Remove DHCP client leases
dhcp-server-state * Remove DHCP server leases
dovecot-data * Remove Dovecot (mail server) data
firewall-rules Remove the firewall rules
flag-reconfiguration Flag the system for reconfiguration
fs-uuids Change filesystem UUIDs
kerberos-data Remove Kerberos data in the guest
logfiles * Remove many log files from the guest
lvm-uuids * Change LVM2 PV and VG UUIDs
machine-id * Remove the local machine ID
mail-spool * Remove email from the local mail spool directory
net-hostname * Remove HOSTNAME in network interface configuration
net-hwaddr * Remove HWADDR (hard-coded MAC address) configuration
pacct-log * Remove the process accounting log files
package-manager-cache * Remove package manager cache
pam-data * Remove the PAM data in the guest
puppet-data-log * Remove the data and log files of puppet
rh-subscription-manager * Remove the RH subscription manager files
rhn-systemid * Remove the RHN system ID
rpm-db * Remove host-specific RPM database files
samba-db-log * Remove the database and log files of Samba
script * Run arbitrary scripts against the guest
smolt-uuid * Remove the Smolt hardware UUID
ssh-hostkeys * Remove the SSH host keys in the guest
ssh-userdir * Remove ".ssh" directories in the guest
sssd-db-log * Remove the database and log files of sssd
tmp-files * Remove temporary files
udev-persistent-net * Remove udev persistent net rules
user-account Remove the user accounts in the guest
utmp * Remove the utmp file
yum-uuid * Remove the yum UUID
```

APPENDIX C. LEGAL NOTICE

Copyright © 2022 Red Hat, Inc.

Licensed under the ([Creative Commons Attribution–ShareAlike 4.0 International License](#)). Derived from documentation for the ([oVirt Project](#)). If you distribute this document or an adaptation of it, you must provide the URL for the original version.

Modified versions must remove all Red Hat trademarks.

Red Hat, Red Hat Enterprise Linux, the Red Hat logo, 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.