



Red Hat Virtualization 4.1

Installation Guide

Installing Red Hat Virtualization

Red Hat Virtualization 4.1 Installation Guide

Installing Red Hat Virtualization

Red Hat Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@redhat.com

Legal Notice

Copyright © 2016 Red Hat.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

A comprehensive guide to installing Red Hat Virtualization.

Table of Contents

CHAPTER 1. INTRODUCTION	3
1.1. SYSTEM REQUIREMENTS	3
PART I. INSTALLING THE RED HAT VIRTUALIZATION MANAGER	4
CHAPTER 2. RED HAT VIRTUALIZATION MANAGER	5
2.1. SUBSCRIBING TO THE REQUIRED ENTITLEMENTS	5
2.2. INSTALLING THE RED HAT VIRTUALIZATION MANAGER PACKAGES	6
2.3. CONFIGURING THE RED HAT VIRTUALIZATION MANAGER	6
2.4. CONNECTING TO THE ADMINISTRATION PORTAL	12
CHAPTER 3. RED HAT VIRTUALIZATION MANAGER RELATED TASKS	14
3.1. REMOVING THE RED HAT VIRTUALIZATION MANAGER	14
3.2. CONFIGURING A LOCAL REPOSITORY FOR OFFLINE RED HAT VIRTUALIZATION MANAGER INSTALLATION	15
PART II. INSTALLING HOSTS	18
CHAPTER 4. INTRODUCTION TO HOSTS	19
4.1. HOST COMPATIBILITY	19
CHAPTER 5. RED HAT VIRTUALIZATION HOSTS	20
5.1. INSTALLING RED HAT VIRTUALIZATION HOST	20
5.2. ADVANCED INSTALLATION	22
CHAPTER 6. RED HAT ENTERPRISE LINUX HOSTS	28
6.1. INSTALLING RED HAT ENTERPRISE LINUX HOSTS	28
6.2. SUBSCRIBING TO THE REQUIRED ENTITLEMENTS	28
6.3. INSTALLING COCKPIT ON RED HAT ENTERPRISE LINUX HOSTS	29
CHAPTER 7. ADDING A HOST TO THE RED HAT VIRTUALIZATION MANAGER	31
PART III. ATTACHING STORAGE	32
CHAPTER 8. STORAGE	33
8.1. INTRODUCTION TO STORAGE	33
8.2. ADDING FCP STORAGE	33
APPENDIX A. CHANGING THE PERMISSIONS FOR THE LOCAL ISO DOMAIN	36
APPENDIX B. ATTACHING THE LOCAL ISO DOMAIN TO A DATA CENTER	37
APPENDIX C. ENABLING GLUSTER PROCESSES ON RED HAT GLUSTER STORAGE NODES	38
APPENDIX D. PREPARING A REMOTE POSTGRESQL DATABASE	39
APPENDIX E. PREPARING A LOCAL MANUALLY-CONFIGURED POSTGRESQL DATABASE FOR USE WITH THE RED HAT VIRTUALIZATION MANAGER	42
APPENDIX F. INSTALLING A WEBSOCKET PROXY ON A SEPARATE MACHINE	44
APPENDIX G. CONFIGURING A HOST FOR PCI PASSTHROUGH	47

CHAPTER 1. INTRODUCTION

This guide covers:

- Installing and configuring the Red Hat Virtualization Manager.
- Installing and configuring hosts.
- Attaching existing FCP storage to your Red Hat Virtualization environment. More storage options can be found in the [Administration Guide](#).



IMPORTANT

To avoid potential timing or authentication issues, configure the Network Time Protocol (NTP) on the hosts, Manager, and other servers in the environment to synchronize with the same NTP server. See [Configuring NTP Using the chrony Suite](#) and [Synchronizing the System Clock with a Remote Server](#) in the *Red Hat Enterprise Linux 7 System Administrator's Guide*.

1.1. SYSTEM REQUIREMENTS

For information about the requirements of a Red Hat Virtualization system, see [Requirements](#) in the *Planning and Prerequisites Guide*.

PART I. INSTALLING THE RED HAT VIRTUALIZATION MANAGER

CHAPTER 2. RED HAT VIRTUALIZATION MANAGER

2.1. SUBSCRIBING TO THE REQUIRED ENTITLEMENTS

Once you have installed a Red Hat Enterprise Linux base operating system and made sure the system meets the requirements listed in the previous chapter, you must register the system with Red Hat Subscription Manager, and subscribe to the required entitlements to install the Red Hat Virtualization Manager packages.

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

```
# subscription-manager register
```

2. Find the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** subscription pools and note down the pool IDs.

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

3. Use the pool IDs located in the previous step to attach the entitlements to the system:

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



NOTE

To find out what subscriptions are currently attached, run:

```
# subscription-manager list --consumed
```

To list all enabled repositories, run:

```
# yum repolist
```

4. Disable all existing repositories:

```
# subscription-manager repos --disable=*
```

5. Enable the required repositories:

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-supplementary-rpms
# subscription-manager repos --enable=rhel-7-server-rhv-4.1-rpms
# subscription-manager repos --enable=rhel-7-server-rhv-4-tools-rpms
# subscription-manager repos --enable=jb-eap-7.1-for-rhel-7-server-rpms
```

You have now subscribed your system to the required entitlements. Proceed to the next section to install the Red Hat Virtualization Manager packages.

2.2. INSTALLING THE RED HAT VIRTUALIZATION MANAGER PACKAGES

Before you can configure and use the Red Hat Virtualization Manager, you must install the `rhev` package and dependencies.

Procedure 2.1. Installing the Red Hat Virtualization Manager Packages

1. To ensure all packages are up to date, run the following command on the machine where you are installing the Red Hat Virtualization Manager:

```
# yum update
```



NOTE

Reboot the machine if any kernel related packages have been updated.

2. Run the following command to install the `rhev` package and dependencies.

```
# yum install rhevm
```

Proceed to the next step to configure your Red Hat Virtualization Manager.

2.3. CONFIGURING THE RED HAT VIRTUALIZATION MANAGER

After you have installed the `rhev` package and dependencies, you must configure the Red Hat Virtualization Manager using the `engine-setup` command. This command asks you a series of questions and, after you provide the required values for all questions, applies that configuration and starts the `ovirt-engine` service.

By default, `engine-setup` creates and configures the Manager database locally on the Manager machine. Alternatively, you can configure the Manager to use a remote database or a manually-configured local database; however, you must set up that database before running `engine-setup`. To set up a remote database see [Appendix D, *Preparing a Remote PostgreSQL Database*](#) . To set up a manually-configured local database, see [Appendix E, *Preparing a Local Manually-Configured PostgreSQL Database for Use with the Red Hat Virtualization Manager*](#).

By default, `engine-setup` will configure a websocket proxy on the Manager. However, for security and performance reasons, the user can choose to configure it on a separate host. See [Appendix F, *Installing a Websocket Proxy on a Separate Machine*](#) for instructions.



IMPORTANT

The **engine-setup** command guides you through several distinct configuration stages, each comprising several steps that require user input. Suggested configuration defaults are provided in square brackets; if the suggested value is acceptable for a given step, press **Enter** to accept that value.

You can run **engine-setup --accept-defaults** to automatically accept all questions that have default answers. This option should be used with caution and only if you are familiar with engine-setup.

Procedure 2.2. Configuring the Red Hat Virtualization Manager

1. Run the **engine-setup** command to begin configuration of the Red Hat Virtualization Manager:

```
# engine-setup
```

2. Press **Enter** to configure the Manager:

```
Configure Engine on this host (Yes, No) [Yes]:
```

3. Optionally allow **engine-setup** to configure the Image I/O Proxy (virt-imageio-proxy) to allow the Manager to upload virtual disks into storage domains. See [Uploading a Disk Image to a Storage Domain](#) in the *Administration Guide* for more information.

```
Configure Image I/O Proxy on this host? (Yes, No) [Yes]:
```

4. Optionally allow **engine-setup** to configure a websocket proxy server for allowing users to connect to virtual machines via the noVNC or HTML 5 consoles:

```
Configure WebSocket Proxy on this machine? (Yes, No) [Yes]:
```

To configure the websocket proxy on a separate machine, select **No** and refer to [Appendix F, Installing a WebSocket Proxy on a Separate Machine](#) for configuration instructions.

5. Choose whether to configure Data Warehouse on the Manager machine.

```
Please note: Data Warehouse is required for the engine. If you
choose to not configure it on this host, you have to configure it on
a remote host, and then configure the engine on this host so that it
can access the database of the remote Data Warehouse host.
Configure Data Warehouse on this host (Yes, No) [Yes]:
```

To configure Data Warehouse on a separate machine, select **No** and see [Installing and Configuring Data Warehouse on a Separate Machine](#) in the *Data Warehouse Guide* for installation and configuration instructions.

6. Optionally allow access to a virtual machines's serial console from the command line.

```
Configure VM Console Proxy on this host (Yes, No) [Yes]:
```

Additional configuration is required on the client machine to use this feature. See [Opening a Serial Console to a Virtual Machine](#) in the *Virtual Machine Management Guide*.

7. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**. Note that the automatically detected hostname may be incorrect if you are using virtual hosts.

```
Host fully qualified DNS name of this server [autodetected host name]:
```

8. The **engine-setup** command checks your firewall configuration and offers to modify that configuration to open the ports used by the Manager for external communication such as TCP ports 80 and 443. If you do not allow **engine-setup** to modify your firewall configuration, then you must manually open the ports used by the Manager.

```
Setup can automatically configure the firewall on this system.  
Note: automatic configuration of the firewall may overwrite current settings.  
Do you want Setup to configure the firewall? (Yes, No) [Yes]:
```

If you choose to automatically configure the firewall, and no firewall managers are active, you are prompted to select your chosen firewall manager from a list of supported options. Type the name of the firewall manager and press **Enter**. This applies even in cases where only one option is listed.

9. Choose to use either a local or remote PostgreSQL database as the Data Warehouse database:

```
Where is the DWH database located? (Local, Remote) [Local]:
```

- If you select **Local**, the **engine-setup** command can configure your database automatically (including adding a user and a database), or it can connect to a preconfigured local database:

```
Setup can configure the local postgresql server automatically for the DWH to run. This may conflict with existing applications.  
Would you like Setup to automatically configure postgresql and create DWH database, or prefer to perform that manually?  
(Automatic, Manual) [Automatic]:
```

- a. If you select **Automatic** by pressing **Enter**, no further action is required here.
- b. If you select **Manual**, input the following values for the manually-configured local database:

```
DWH database secured connection (Yes, No) [No]:  
DWH database name [ovirt_engine_history]:  
DWH database user [ovirt_engine_history]:  
DWH database password:
```

**NOTE**

engine-setup requests these values after the Manager database is configured in the next step.

- If you select **Remote**, input the following values for the preconfigured remote database host:

```
DWH database host [localhost]:
DWH database port [5432]:
DWH database secured connection (Yes, No) [No]:
DWH database name [ovirt_engine_history]:
DWH database user [ovirt_engine_history]:
DWH database password:
```

**NOTE**

engine-setup requests these values after the Manager database is configured in the next step.

10. Choose to use either a local or remote PostgreSQL database as the Manager database:

```
Where is the Engine database located? (Local, Remote) [Local]:
```

- If you select **Local**, the **engine-setup** command can configure your database automatically (including adding a user and a database), or it can connect to a preconfigured local database:

```
Setup can configure the local postgresql server automatically for
the engine to run. This may conflict with existing applications.
Would you like Setup to automatically configure postgresql and
create Engine database, or prefer to perform that manually?
(Automatic, Manual) [Automatic]:
```

- a. If you select **Automatic** by pressing **Enter**, no further action is required here.
- b. If you select **Manual**, input the following values for the manually-configured local database:

```
Engine database secured connection (Yes, No) [No]:
Engine database name [engine]:
Engine database user [engine]:
Engine database password:
```

- If you select **Remote**, input the following values for the preconfigured remote database host:

```
Engine database host [localhost]:
Engine database port [5432]:
Engine database secured connection (Yes, No) [No]:
```

```
Engine database name [engine]:  
Engine database user [engine]:  
Engine database password:
```

11. Set a password for the automatically created administrative user of the Red Hat Virtualization Manager:

```
Engine admin password:  
Confirm engine admin password:
```

12. Select **Gluster**, **Virt**, or **Both**:

```
Application mode (Both, Virt, Gluster) [Both]:
```

Both offers the greatest flexibility. In most cases, select **Both**. **Virt** application mode allows you to run virtual machines in the environment; **Gluster** application mode only allows you to manage GlusterFS from the Administration Portal.

13. Set the default value for the **wipe_after_delete** flag, which wipes the blocks of a virtual disk when the disk is deleted.

```
Default SAN wipe after delete (Yes, No) [No]:
```

14. The Manager uses certificates to communicate securely with its hosts. This certificate can also optionally be used to secure HTTPS communications with the Manager. Provide the organization name for the certificate:

```
Organization name for certificate [autodetected domain-based name]:
```

15. Optionally allow **engine-setup** to make the landing page of the Manager the default page presented by the Apache web server:

```
Setup can configure the default page of the web server to present  
the application home page. This may conflict with existing  
applications.
```

```
Do you wish to set the application as the default web page of the  
server? (Yes, No) [Yes]:
```

16. By default, external SSL (HTTPS) communication with the Manager is secured with the self-signed certificate created earlier in the configuration to securely communicate with hosts. Alternatively, choose another certificate for external HTTPS connections; this does not affect how the Manager communicates with hosts:

```
Setup can configure apache to use SSL using a certificate issued  
from the internal CA.
```

```
Do you wish Setup to configure that, or prefer to perform that  
manually? (Automatic, Manual) [Automatic]:
```

17. Choose whether or not to create an NFS share on the Manager to use as an ISO storage domain. The local ISO domain provides a selection of images that can be used in the initial setup of virtual machines:

Configure an NFS share on this server to be used as an ISO Domain? (Yes, No) [No]:

- If you select the default (**No**), no further action is required here.
- If you select **Yes**, then you need to provide extra information to set up the ISO domain:
 - a. Specify the path for the ISO domain:

Local ISO domain path [/var/lib/exports/iso]:

- b. Specify the networks or hosts that require access to the ISO domain:

Local ISO domain ACL: *10.1.2.0/255.255.255.0(rw)*
host01.example.com(rw) host02.example.com(rw)

The example above allows access to a single /24 network and two specific hosts. See the **exports(5)** man page for further formatting options.

- c. Specify a display name for the ISO domain:

Local ISO domain name [ISO_DOMAIN]:

18. Choose how long Data Warehouse will retain collected data:



NOTE

This step is skipped if you chose not to configure Data Warehouse on the Manager machine.

Please choose Data Warehouse sampling scale:

- (1) Basic
 - (2) Full
- (1, 2)[1]:

Full uses the default values for the data storage settings listed in the [Data Warehouse Guide](#) (recommended when Data Warehouse is installed on a remote host).

Basic reduces the values of **DWH_TABLES_KEEP_HOURLY** to **720** and **DWH_TABLES_KEEP_DAILY** to **0**, easing the load on the Manager machine (recommended when the Manager and Data Warehouse are installed on the same machine).

19. Review the installation settings, and press **Enter** to accept the values and proceed with the installation:

Please confirm installation settings (OK, Cancel) [OK]:

When your environment has been configured, **engine-setup** displays details about how to access your environment. If you chose to manually configure the firewall, **engine-setup** provides a custom list of ports that need to be opened, based on the

options selected during setup. The **engine-setup** command also saves your answers to a file that can be used to reconfigure the Manager using the same values, and outputs the location of the log file for the Red Hat Virtualization Manager configuration process.

20. If you intend to link your Red Hat Virtualization environment with a directory server, configure the date and time to synchronize with the system clock used by the directory server to avoid unexpected account expiry issues. See [Synchronizing the System Clock with a Remote Server](#) in the *Red Hat Enterprise Linux System Administrator's Guide* for more information.
21. Install the certificate authority according to the instructions provided by your browser. You can get the certificate authority's certificate by navigating to **`http://your-manager-fqdn/ovirt-engine/services/pki-resource?resource=ca-certificate&format=X509-PEM-CA`**, replacing *your-manager-fqdn* with the fully qualified domain name (FQDN) that you provided during the installation.

Proceed to the next section to connect to the Administration Portal as the **admin@internal** user. Then, proceed with setting up hosts, and attaching storage.

2.4. CONNECTING TO THE ADMINISTRATION PORTAL

Access the Administration Portal using a web browser.

1. In a web browser, navigate to **`https://your-manager-fqdn/ovirt-engine`**, replacing *your-manager-fqdn* with the fully qualified domain name that you provided during installation.



NOTE


You can access the Administration Portal using alternate host names or IP addresses. To do so, you need to add a configuration file under **`/etc/ovirt-engine/engine.conf.d/`**. For example:

```
# vi /etc/ovirt-engine/engine.conf.d/99-custom-sso-
setup.conf
SSO_ALTERNATE_ENGINE_FQDNS="alias1.example.com
alias2.example.com"
```

The list of alternate host names needs to be separated by spaces. You can also add the IP address of the Manager to the list, but using IP addresses instead of DNS-resolvable host names is not recommended.

2. Click **Administration Portal**. An SSO login page displays. SSO login enables you to log in to the Administration and User Portals at the same time.
3. Enter your **User Name** and **Password**. If you are logging in for the first time, use the user name **admin** in conjunction with the password that you specified during installation.
4. Select the domain against which to authenticate from the **Domain** list. If you are logging in using the internal **admin** user name, select the **internal** domain.
5. Click **Log In**.

6. You can view the Administration Portal in multiple languages. The default selection will be chosen based on the locale settings of your web browser. If you would like to view the Administration Portal in a language other than the default, select your preferred language from the drop-down list on the welcome page.



IMPORTANT

Keep the environment up-to-date. See <https://access.redhat.com/articles/2974891> for more information. Since bug fixes for known issues are frequently released, Red Hat recommends using scheduled tasks to update the hosts and the Manager.

To log out of the Red Hat Virtualization Administration Portal, click your user name in the header bar and click **Sign Out**. You are logged out of all portals and the Manager welcome screen displays.

The next chapter contains additional Manager related tasks which are optional. If the tasks are not applicable to your environment, proceed to [Part II, “Installing Hosts”](#).

CHAPTER 3. RED HAT VIRTUALIZATION MANAGER RELATED TASKS

3.1. REMOVING THE RED HAT VIRTUALIZATION MANAGER

You can use the **engine-cleanup** command to remove specific components or all components of the Red Hat Virtualization Manager.



NOTE

A backup of the engine database and a compressed archive of the PKI keys and configuration are always automatically created. These files are saved under `/var/lib/ovirt-engine/backups/`, and include the date and **engine-** and **engine-pki-** in their file names respectively.

Procedure 3.1. Removing the Red Hat Virtualization Manager

1. Run the following command on the machine on which the Red Hat Virtualization Manager is installed:

```
# engine-cleanup
```

2. You are prompted whether to remove all Red Hat Virtualization Manager components:

- Type **Yes** and press **Enter** to remove all components:

```
Do you want to remove all components? (Yes, No) [Yes]:
```

- Type **No** and press **Enter** to select the components to remove. You can select whether to retain or remove each component individually:

```
Do you want to remove Engine database content? All data will be lost (Yes, No) [No]:
```

```
Do you want to remove PKI keys? (Yes, No) [No]:
```

```
Do you want to remove PKI configuration? (Yes, No) [No]:
```

```
Do you want to remove Apache SSL configuration? (Yes, No) [No]:
```

3. You are given another opportunity to change your mind and cancel the removal of the Red Hat Virtualization Manager. If you choose to proceed, the **ovirt-engine** service is stopped, and your environment's configuration is removed in accordance with the options you selected.

```
During execution engine service will be stopped (OK, Cancel) [OK]:
ovirt-engine is about to be removed, data will be lost (OK, Cancel)
[Cancel]:OK
```

4. Remove the Red Hat Virtualization packages:

```
# yum remove rhevm* vds-m-bootstrap
```

3.2. CONFIGURING A LOCAL REPOSITORY FOR OFFLINE RED HAT VIRTUALIZATION MANAGER INSTALLATION

To install Red Hat Virtualization Manager on a system that does not have a direct connection to the Content Delivery Network, download the required packages on a system that has Internet access, then create a repository that can be shared with the offline Manager machine. The system hosting the repository must be connected to the same network as the client systems where the packages are to be installed.

1. Install Red Hat Enterprise Linux 7 Server on a system that has access to the Content Delivery Network. This system downloads all the required packages, and distributes them to your offline system(s).



IMPORTANT

Ensure that the system used in this procedure has a large amount of free disk space available. This procedure downloads a large number of packages, and requires up to 50GB of free disk space.

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

```
# subscription-manager register
```

3. Subscribe the system to all required entitlements:

1. Find the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** subscription pools and note down the pool IDs.

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

2. Use the pool IDs located in the previous step to attach the entitlements to the system:

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

3. Disable all existing repositories:

```
# subscription-manager repos --disable=*
```

4. Enable the required repositories:

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-
supplementary-rpms
# subscription-manager repos --enable=rhel-7-server-rhv-4.1-rpms
# subscription-manager repos --enable=rhel-7-server-rhv-4-tools-
rpms
# subscription-manager repos --enable=jb-eap-7.1-for-rhel-7-
server-rpms
```

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

```
# yum update
```

**NOTE**

Reboot the machine if any kernel related packages have been updated.

4. Servers that are not connected to the Internet can access software repositories on other systems using File Transfer Protocol (FTP). To create the FTP repository, install and configure vsftpd:

- a. Install the vsftpd package:

```
# yum install vsftpd
```

- b. Start the **vsftpd** service, and ensure the service starts on boot:

```
# systemctl start vsftpd.service  
# systemctl enable vsftpd.service
```

- c. Create a sub-directory inside the **/var/ftp/pub/** directory. This is where the downloaded packages will be made available:

```
# mkdir /var/ftp/pub/rhevrepo
```

5. Download packages from all configured software repositories to the **rhevrepo** directory. This includes repositories for all Content Delivery Network subscription pools the system is subscribed to, and any locally configured repositories:

```
# reposync -l -p /var/ftp/pub/rhevrepo
```

This command downloads a large number of packages, and takes a long time to complete. The **-l** option enables yum plug-in support.

6. Install the createrepo package:

```
# yum install createrepo
```

7. Create repository metadata for each of the sub-directories where packages were downloaded under **/var/ftp/pub/rhevrepo**:

```
# for DIR in `find /var/ftp/pub/rhevrepo -maxdepth 1 -mindepth 1 -  
type d`; do createrepo $DIR; done;
```

8. Create a repository file, and copy it to the **/etc/yum.repos.d/** directory on the offline machine on which you will install the Manager.

The configuration file can be created manually or with a script. Run the script below on the system hosting the repository, replacing **ADDRESS** in the **baseurl** with the IP address or fully qualified domain name of the system hosting the repository:

```
#!/bin/sh

REPOFILE="/etc/yum.repos.d/rhev.repo"
echo -e " " > $REPOFILE

for DIR in `find /var/ftp/pub/rhevrepo -maxdepth 1 -mindepth 1 -type
d`;
do
    echo -e "[`basename $DIR`]" >> $REPOFILE
    echo -e "name=`basename $DIR`" >> $REPOFILE
    echo -e "baseurl=ftp://ADDRESS/pub/rhevrepo/`basename $DIR`" >>
$REPOFILE
    echo -e "enabled=1" >> $REPOFILE
    echo -e "gpgcheck=0" >> $REPOFILE
    echo -e "\n" >> $REPOFILE
done;
```

9. Install the Manager packages on the offline system. See [Section 2.2, “Installing the Red Hat Virtualization Manager Packages”](#) for instructions. Packages are installed from the local repository, instead of from the Content Delivery Network.
10. Configure the Manager. See [Section 2.3, “Configuring the Red Hat Virtualization Manager”](#) for initial configuration instructions.
11. Continue with host, storage, and virtual machine configuration.

PART II. INSTALLING HOSTS

CHAPTER 4. INTRODUCTION TO HOSTS

Red Hat Virtualization supports two types of hosts: Red Hat Virtualization Host (RHVH) and Red Hat Enterprise Linux host. Depending on your environment requirement, you may want to use one type only or both in your Red Hat Virtualization environment. It is recommended that you install and attach at least two hosts to the Red Hat Virtualization environment. Where you attach only one host you will be unable to access features such as migration and high availability.



IMPORTANT

SELinux is in enforcing mode upon installation. To verify, run **getenforce**. SELinux is required to be in enforcing mode on all hypervisors and Managers for your Red Hat Virtualization environment to be supported by Red Hat.

Table 4.1. Hosts

Host Type	Other Names	Description
Red Hat Virtualization Host	RHVH, thin host	This is a minimal operating system based on Red Hat Enterprise Linux. It is distributed as an ISO file from the Customer Portal and contains only the packages required for the machine to act as a host.
Red Hat Enterprise Linux Host	RHEL-based hypervisor, thick host	Red Hat Enterprise Linux hosts subscribed to the appropriate entitlements can be used as hosts.

4.1. HOST COMPATIBILITY

When you create a new data center, you can set the compatibility version. Select the compatibility version that suits all the hosts in the data center. Once set, version regression is not allowed. For a fresh Red Hat Virtualization installation, the latest compatibility version is set in the default data center and default cluster; to use an earlier compatibility version, you must create additional data centers and clusters. For more information about compatibility versions see *Red Hat Virtualization Manager Compatibility* in the [Red Hat Virtualization Life Cycle](#).

CHAPTER 5. RED HAT VIRTUALIZATION HOSTS

5.1. INSTALLING RED HAT VIRTUALIZATION HOST

Red Hat Virtualization Host (RHVH) is a minimal operating system based on Red Hat Enterprise Linux that is designed to provide a simple method for setting up a physical machine to act as a hypervisor in a Red Hat Virtualization environment. The minimal operating system contains only the packages required for the machine to act as a hypervisor, and features a Cockpit user interface for monitoring the host and performing administrative tasks. See <http://cockpit-project.org/running.html> for the minimum browser requirements.



NOTE

Importing RHVH images into Satellite from configured repositories is not supported.

RHVH supports NIST 800-53 partitioning requirements to improve security. RHVH uses a NIST 800-53 partition layout by default.

Before you proceed, ensure the machine on which you are installing RHVH meets the hardware requirements listed in [Host Requirements](#) in the *Planning and Prerequisites Guide*.

Installing RHVH on a physical machine involves three key steps:

1. Download the RHVH ISO image from the Customer Portal.
2. Write the RHVH ISO image to a USB, CD, or DVD.
3. Install the RHVH minimal operating system.


Procedure 5.1. Installing Red Hat Virtualization Host

1. Download the RHVH ISO image from the Customer Portal:
 - a. Log in to the Customer Portal at <https://access.redhat.com>.
 - b. Click **Downloads** in the menu bar.
 - c. Click **Red Hat Virtualization**, scroll up, and click **Download Latest** to access the product download page.
 - d. Choose the appropriate hypervisor image and click **Download Now**.
 - e. Create a bootable media device. See [Making Media](#) in the *Red Hat Enterprise Linux Installation Guide* for more information.
2. Start the machine on which to install RHVH using the prepared installation media.
3. From the boot menu, select the **Install** option, and press **Enter**.

**NOTE**

You can also press the **Tab** key to edit the kernel parameters. Kernel parameters must be separated by a space, and you can boot the system using the specified kernel parameters by pressing the **Enter** key. Press the **Esc** key to clear any changes to the kernel parameters and return to the boot menu.

4. Select a language, and click **Continue**.
5. Select a time zone from the **Date & Time** screen and click **Done**.
6. Select a keyboard layout from the **Keyboard** screen and click **Done**.
7. Select the device on which to install RHVH from the **Installation Destination** screen. Optionally, enable encryption. Click **Done**.

**IMPORTANT**

Red Hat strongly recommends using the **Automatically configure partitioning** option. However, if you do select **I will configure partitioning**, see [Section 5.2.1, “Custom Partitioning”](#) for details.

**NOTE**

For information on preserving local storage domains when reinstalling RHVH, see [Upgrading to RHVH While Preserving Local Storage](#) in the *Upgrade Guide* for Red Hat Virtualization 4.0 for more details.

8. Select a network from the **Network & Host Name** screen and click **Configure...** to configure the connection details. Enter a host name in the **Host name** field, and click **Done**.
9. Optionally configure **Language Support**, **Security Policy**, and **Kdump**. See [Installing Using Anaconda](#) in the *Red Hat Enterprise Linux 7 Installation Guide* for more information on each of the sections in the **Installation Summary** screen.
10. Click **Begin Installation**.
11. Set a root password and, optionally, create an additional user while RHVH installs.

**WARNING**

Red Hat strongly recommends not creating untrusted users on RHVH, as this can lead to exploitation of local security vulnerabilities.

12. Click **Reboot** to complete the installation.

**NOTE**

When RHVH restarts, **nodectl check** performs a health check on the host and displays the result when you log in on the command line. The message **node status: OK** or **node status: DEGRADED** indicates the health status. Run **nodectl check** to get more information. The service is enabled by default.

13. Once the installation is complete, log in to the Cockpit user interface at `https://HostFQDNorIP:9090` to subscribe the host to the Content Delivery Network. Navigate to the **Subscriptions** sub-tab, click **Register System**, and enter your Customer Portal username and password. The system automatically subscribes to the **Red Hat Virtualization Host** entitlement.
14. Click **Terminal**, and enable the **Red Hat Virtualization Host 7** repository to allow later updates to the Red Hat Virtualization Host:

```
# subscription-manager repos --enable=rhel-7-server-rhvh-4-rpms
```

You can now add the host to your Red Hat Virtualization environment. See [Chapter 7, Adding a Host to the Red Hat Virtualization Manager](#)


5.2. ADVANCED INSTALLATION

5.2.1. Custom Partitioning

Custom partitioning on Red Hat Virtualization Host (RHVH) is not recommended. Red Hat strongly recommends using the **Automatically configure partitioning** option in the **Installation Destination** window.

If your installation requires custom partitioning, select the **I will configure partitioning** option during the installation, and note that the following restrictions apply:

- You must select the **LVM Thin Provisioning** option in the **Manual Partitioning** window.
- The following directories are required and must be on thin provisioned logical volumes:
 - `root (/)`
 - `/home`
 - `/tmp`
 - `/var`
 - `/var/log`
 - `/var/log/audit`



IMPORTANT

Do not create a separate partition for `/usr`. Doing so will cause the installation to fail.

`/usr` must be on a logical volume that is able to change versions along with RHVH, and therefore should be left on root (`/`).

For information about the required storage sizes for each partition, see [Storage Requirements](#) in the *Planning and Prerequisites Guide*.

- The `/boot` directory should be defined as a standard partition.
- The `/var` directory must be on a separate volume or disk.
- Only XFS or Ext4 file systems are supported.

Example 5.1. Configuring Manual Partitioning in a Kickstart File

The following example demonstrates how to configure manual partitioning in a Kickstart file. See [Section 5.2.2, “Automating Red Hat Virtualization Host Deployment”](#) for more information on RHVH Kickstart files.

```
clearpart --all
part /boot --fstype xfs --size=1000 --ondisk=sda
part pv.01 --size=42000 --grow
volgroup HostVG pv.01 --reserved-percent=20
logvol swap --vgname=HostVG --name=swap --fstype=swap --recommended
logvol none --vgname=HostVG --name=HostPool --thinpool --size=40000 --
grow
logvol / --vgname=HostVG --name=root --thin --fstype=ext4 --
poolname=HostPool --fsoptions="defaults,discard" --size=6000 --grow
logvol /var --vgname=HostVG --name=var --thin --fstype=ext4 --
poolname=HostPool
--fsoptions="defaults,discard" --size=15000
logvol /var/log --vgname=HostVG --name=var_log --thin --fstype=ext4 --
poolname=HostPool --fsoptions="defaults,discard" --size=8000
logvol /var/log/audit --vgname=HostVG --name=var_audit --thin --
fstype=ext4 --poolname=HostPool --fsoptions="defaults,discard" --
size=2000
logvol /home --vgname=HostVG --name=home --thin --fstype=ext4 --
poolname=HostPool --fsoptions="defaults,discard" --size=1000
logvol /tmp --vgname=HostVG --name=tmp --thin --fstype=ext4 --
poolname=HostPool --fsoptions="defaults,discard" --size=1000
```



NOTE

If you use `logvol --thinpool --grow`, you must also include `volgroup --reserved-space` or `volgroup --reserved-percent` to reserve space in the volume group for the thin pool to grow.

5.2.2. Automating Red Hat Virtualization Host Deployment

You can install Red Hat Virtualization Host (RHVH) without a physical media device by booting from the network using PXE. You can automate the installation process by using a Kickstart file containing the answers to the installation questions. The Kickstart file can also be accessed over the network, removing the need for physical media.

Instructions for both tasks can be found in the [Red Hat Enterprise Linux 7 Installation Guide](#), as RHVH is installed in much the same way as Red Hat Enterprise Linux. The main differences required for RHVH are included in the following procedure.

This procedure also includes example files for use with Red Hat Satellite.

Procedure 5.2. Automating Deployment using PXE and Kickstart

1. Download the RHVH ISO image from the Customer Portal:
 - a. Log in to the Customer Portal at <https://access.redhat.com>.
 - b. Click **Downloads** in the menu bar.
 - c. Click **Red Hat Virtualization**, scroll up, and click **Download Latest** to access the product download page.
 - d. Choose the appropriate hypervisor image and click **Download Now**.
2. Make the RHVH ISO image available over the network using the instructions in [Installation Source on a Network](#).
3. Extract the **squashfs.img** file from the RHVH ISO:

```
# mount -o loop /path/to/RHVH-ISO /mnt/rhvh
# cp /mnt/rhvh/Packages/redhat-virtualization-host-image-update*
/tmp
# cd /tmp
# rpm2cpio redhat-virtualization-host-image-update* | cpio -idmv
```

The **squashfs.img** file is located at **/tmp/usr/share/redhat-virtualization-host/image/**.

4. Configure the PXE server using the instructions in [Preparing for a Network Installation](#).

The following requirements apply in order to boot RHVH from the PXE server:

- Ensure that you copy the RHVH boot images to the **/tftpboot** directory.

```
# cp URL/to/RHVH-ISO/images/pxeboot/{vmlinuz,initrd.img}
/var/lib/tftpboot/pxelinux/
```

- The boot loader configuration file must include a RHVH label that specifies the RHVH boot images.

```
LABEL rhvh
MENU LABEL Install Red Hat Virtualization ^Host
KERNEL /var/lib/tftpboot/pxelinux/vmlinuz
APPEND initrd=/var/lib/tftpboot/pxelinux/initrd.img
inst.stage2=URL/to/RHVH-ISO
```

Example 5.2. Red Hat Virtualization Host PXELinux Satellite Template

The following is an example of a boot loader label that uses information from Red Hat Satellite to provision the host. You must create a global or host group level parameter called **rhvh_image** and populate it with the directory URL where the ISO is mounted or extracted.

```
<%#
kind: PXELinux
name: RHVH PXELinux
%>
# Created for booting new hosts
#

DEFAULT rhvh

LABEL rhvh
KERNEL <%= @kernel %>
APPEND initrd=<%= @initrd %> inst.ks=<%= foreman_url("provision")
%> inst.stage2=<%= @host.params["rhvh_image"] %> intel_iommu=on
console=tty0 console=ttyS1,115200n8 ssh_pwauth=1
local_boot_trigger=<%= foreman_url("built") %>
IPAPPEND 2
```

5. Create a Kickstart file and make it available over the network using the instructions in [Kickstart Installations](#).

The following constraints apply to RHVH Kickstart files:

- The **%packages** section is not required for RHVH. Instead, use the **liveimg** option and specify the **squashfs.img** file from the RHVH ISO image.

```
liveimg --url=example.com/tmp/usr/share/redhat-virtualization-
host/image/squashfs.img
```

- The **autopart** command is highly recommended. Thin provisioning must be used.

```
autopart --type=thinp
```

**NOTE**

The **--no-home** option does not work in RHVH. This is an expected behavior, because **/home** is a required directory.

If your installation requires manual partitioning instead, see [Section 5.2.1, “Custom Partitioning”](#) for a list of limitations that apply to partitions, and an example of manual partitioning in a Kickstart file.

- A **%post** section that calls the **nodectl init** command is required.

```
%post
nodectl init
%end
```

To fully automate the installation process, you can add this Kickstart file to the boot loader configuration file on the PXE server. Specify the Kickstart location by adding **inst.ks=** to the **APPEND** line:

```
APPEND initrd=/var/tftpboot/pxelinux/initrd.img
inst.stage2=URL/to/RHVH-ISO inst.ks=URL/to/RHVH-ks.cfg
```

Example 5.3. Red Hat Virtualization Host Kickstart File

The following is an example of a Kickstart file used to deploy Red Hat Virtualization Host. You can include additional commands and options as required.

```
liveimg --url=http://1.2.3.4/tmp/usr/share/redhat-virtualization-
host/image/squashfs.img
clearpart --all
autopart --type=thinp
rootpw --plaintext ovirt
timezone --utc America/Phoenix
zerombr
text

reboot

%post --erroronfail
nodectl init
%end
```

Example 5.4. Red Hat Virtualization Host Kickstart File with Registration and Network Configuration

The following is an example of a Kickstart file that uses information from Red Hat Satellite to configure the host network and register the host to the Satellite server. You must create a global or host group level parameter called **rhvh_image** and populate it with the directory URL to the **squashfs.img** file. **ntp_server1** is also a global or host group level variable.

```
<%=#
kind: provision
name: RHVH Kickstart default
oses:
- RHVH
%>
install
liveimg --url=<%= @host.params['rhvh_image'] %>squashfs.img

network --bootproto static --ip=<%= @host.ip %> --netmask=<%=
@host.subnet.mask %> --gateway=<%= @host.subnet.gateway %> --
nameserver=<%= @host.subnet.dns_primary %> --hostname <%=
```

```
@host.name %>

zerombr
clearpart --all
autopart --type=thinp

rootpw --iscrypted <%= root_pass %>

# installation answers
lang en_US.UTF-8
timezone <%= @host.params['time-zone'] || 'UTC' %>
keyboard us
firewall --service=ssh
services --enabled=sshd

text
reboot

%post --log=/root/ks.post.log --erroronfail
nodedctl init
<%= snippet 'subscription_manager_registration' %>
<%= snippet 'kickstart_networking_setup' %>
/usr/sbin/ntpdate -sub <%= @host.params['ntp_server1'] ||
'0.fedora.pool.ntp.org' %>
/usr/sbin/hwclock --systohc

/usr/bin/curl <%= foreman_url('built') %>

sync
systemctl reboot
%end
```

6. Install RHVH using the instructions in [Booting the Installation on AMD64 and Intel 64 Systems from the Network Using PXE](#).

CHAPTER 6. RED HAT ENTERPRISE LINUX HOSTS

6.1. INSTALLING RED HAT ENTERPRISE LINUX HOSTS

A Red Hat Enterprise Linux host, also known as a RHEL-based hypervisor is based on a standard basic installation of Red Hat Enterprise Linux on a physical server, with **Red Hat Enterprise Linux Server** and the **Red Hat Virtualization** entitlements enabled. For detailed installation instructions, see [Red Hat Enterprise Linux 7 Installation Guide](#)

See [Appendix G, Configuring a Host for PCI Passthrough](#) for more information on how to enable the host hardware and software for device passthrough.



IMPORTANT

Virtualization must be enabled in your host's BIOS settings. For information on changing your host's BIOS settings, refer to your host's hardware documentation.



IMPORTANT

Third-party watchdogs should not be installed on Red Hat Enterprise Linux hosts, as they can interfere with the watchdog daemon provided by VDSM.

6.2. SUBSCRIBING TO THE REQUIRED ENTITLEMENTS

To be used as a virtualization host, make sure the Red Hat Enterprise Linux host meets the hardware requirements listed in [Host Requirements](#) in the *Planning and Prerequisites Guide*. The host must also be registered and subscribed to a number of entitlements using Subscription Manager. Follow this procedure to register with the Content Delivery Network and attach the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** entitlements to the host.

Procedure 6.1. Subscribing to Required Entitlements using Subscription Manager

1. Register your system with the Content Delivery Network, entering your Customer Portal **Username** and **Password** when prompted:


```
# subscription-manager register
```

2. Find the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** subscription pools and note down the pool IDs.

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

3. Use the pool IDs located in the previous step to attach the entitlements to the system:

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


**NOTE**

To find out what subscriptions are currently attached, run:

```
# subscription-manager list --consumed
```

To list all enabled repositories, run:

```
# yum repolist
```

4. Disable all existing repositories:

```
# subscription-manager repos --disable=*
```

5. Enable the required repositories:

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-rhv-4-mgmt-
agent-rpms
# subscription-manager repos --enable=rhel-7-server-ansible-2-rpms
```

If you are installing Red Hat Enterprise Linux 7 hosts, little endian on IBM POWER8 hardware, enable the following repositories instead:

```
# subscription-manager repos --enable=rhel-7-server-rhv-4-mgmt-
agent-for-power-le-rpms
# subscription-manager repos --enable=rhel-7-for-power-le-rpms
```

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

```
# yum update
```

**NOTE**

Reboot the machine if any kernel related packages have been updated.

Once you have subscribed the host to the required entitlements, proceed to [Chapter 7, Adding a Host to the Red Hat Virtualization Manager](#) to attach your host to your Red Hat Virtualization environment.

6.3. INSTALLING COCKPIT ON RED HAT ENTERPRISE LINUX HOSTS

You can install a Cockpit user interface for monitoring the host's resources and performing administrative tasks.

Procedure 6.2. Installing Cockpit on Red Hat Enterprise Linux Hosts

1. By default, when adding a host to the Red Hat Virtualization Manager, the Manager configures the required firewall ports. See [Chapter 7, Adding a Host to the Red Hat](#)

[Virtualization Manager](#) for details. However, if you disable **Automatically configure host firewall** when adding the host, manually configure the firewall according to [Host Firewall Requirements](#) in the *Planning and Prerequisites Guide*.

2. Install the dashboard packages:

```
# yum install cockpit-ovirt-dashboard
```

3. Enable and start the cockpit.socket service:

```
# systemctl enable cockpit.socket  
# systemctl start cockpit.socket
```

4. You can log in to the Cockpit user interface at <https://HostFQDNorIP:9090>.


CHAPTER 7. ADDING A HOST TO THE RED HAT VIRTUALIZATION MANAGER

Adding a host to your Red Hat Virtualization environment can take some time, as the following steps are completed by the platform: virtualization checks, installation of packages, creation of bridge, and a reboot of the host. Use the details pane to monitor the process as the host and the Manager establish a connection.

Procedure 7.1. Adding a Host to the Red Hat Virtualization Manager

1. From the Administration Portal, click the **Hosts** resource tab.
2. Click **New**.
3. Use the drop-down list to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name** and the **Address** of the new host. The standard SSH port, port 22, is auto-filled in the **SSH Port** field.
5. Select an authentication method to use for the Manager to access the host.
 - Enter the root user's password to use password authentication.
 - Alternatively, copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the host to use public key authentication.
6. Click the **Advanced Parameters** button to expand the advanced host settings.
 - a. Optionally disable automatic firewall configuration.
 - b. Optionally add a host SSH fingerprint to increase security. You can add it manually, or fetch it automatically.
7. Optionally configure power management, where the host has a supported power management card. For information on power management configuration, see [Host Power Management Settings Explained](#) in the *Administration Guide*.
8. Click **OK**.

The new host displays in the list of hosts with a status of **Installing**, and you can view the progress of the installation in the details pane. After a brief delay the host status changes to **Up**.



IMPORTANT

Keep the environment up-to-date. See <https://access.redhat.com/articles/2974891> for more information. Since bug fixes for known issues are frequently released, Red Hat recommends using scheduled tasks to update the hosts and the Manager.

PART III. ATTACHING STORAGE


CHAPTER 8. STORAGE

8.1. INTRODUCTION TO STORAGE

A storage domain is a collection of images that have a common storage interface. A storage domain contains complete images of templates and virtual machines (including snapshots), ISO files, and metadata about themselves. A storage domain can be made of either block devices (SAN - iSCSI or FCP) or a file system (NAS - NFS, GlusterFS, or other POSIX compliant file systems).

There are three types of storage domain:

- **Data Domain:** A data domain holds the virtual hard disks and OVF files of all the virtual machines and templates in a data center, and cannot be shared across data centers. Data domains of multiple types (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center, provided they are all shared, rather than local, domains.



IMPORTANT

You must have one host with the status of **Up** and have attached a data domain to a data center before you can attach an ISO domain and an export domain.

- **ISO Domain:** ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines, and can be shared across different data centers. An ISO domain removes the data center's need for physical media. ISO domains can only be NFS-based. Only one ISO domain can be added to a data center.
- **Export Domain:** Export domains are temporary storage repositories that are used to copy and move images between data centers and Red Hat Virtualization environments. Export domains can be used to backup virtual machines. An export domain can be moved between data centers, however, it can only be active in one data center at a time. Export domains can only be NFS-based. Only one export domain can be added to a data center.

See the next section to attach existing FCP storage as a data domain. More storage options are available in the [Administration Guide](#)

8.2. ADDING FCP STORAGE

Red Hat Virtualization platform supports SAN storage by creating a storage domain from a volume group made of pre-existing LUNs. Neither volume groups nor LUNs can be attached to more than one storage domain at a time.

Red Hat Virtualization system administrators need a working knowledge of Storage Area Networks (SAN) concepts. SAN usually uses Fibre Channel Protocol (FCP) for traffic between hosts and shared external storage. For this reason, SAN may occasionally be referred to as FCP storage.

For information regarding the setup and configuration of FCP or multipathing on Red Hat Enterprise Linux, see the [Storage Administration Guide](#) and [DM Multipath Guide](#).

The following procedure shows you how to attach existing FCP storage to your Red Hat Virtualization environment as a data domain. For more information on other supported storage types, see [Storage](#) in the *Administration Guide*.


Procedure 8.1. Adding FCP Storage

1. Click the **Storage** resource tab to list all storage domains.
2. Click **New Domain** to open the **New Domain** window.
3. Enter the **Name** of the storage domain.

LUN ID	Dev. Size	Additional Size	#path	Vendor ID	Product ID	Serial
--------	-----------	-----------------	-------	-----------	------------	--------

Figure 8.1. Adding FCP Storage

4. Use the **Data Center** drop-down menu to select an FCP data center.
If you do not yet have an appropriate FCP data center, select **(none)**.
5. Use the drop-down menus to select the **Domain Function** and the **Storage Type**. The storage domain types that are not compatible with the chosen data center are not available.
6. Select an active host in the **Use Host** field. If this is not the first data domain in a data center, you must select the data center's SPM host.



IMPORTANT

All communication to the storage domain is through the selected host and not directly from the Red Hat Virtualization Manager. At least one active host must exist in the system and be attached to the chosen data center. All hosts must have access to the storage device before the storage domain can be configured.

7. The **New Domain** window automatically displays known targets with unused LUNs when **Data / Fibre Channel** is selected as the storage type. Select the **LUN ID** check box to select all of the available LUNs.
8. Optionally, you can configure the advanced parameters.
 - a. Click **Advanced Parameters**.
 - b. Enter a percentage value into the **Warning Low Space Indicator** field. If the free space available on the storage domain is below this percentage, warning messages are displayed to the user and logged.
 - c. Enter a GB value into the **Critical Space Action Blocker** field. If the free space available on the storage domain is below this value, error messages are displayed to the user and logged, and any new action that consumes space, even temporarily, will be blocked.
 - d. Select the **Wipe After Delete** check box to enable the wipe after delete option. This option can be edited after the domain is created, but doing so will not change the wipe after delete property of disks that already exist.
 - e. Select the **Discard After Delete** check box to enable the discard after delete option. This option can be edited after the domain is created. This option is only available to block storage domains.
9. Click **OK** to create the storage domain and close the window.

The new FCP data domain displays on the **Storage** tab. It will remain with a **Locked** status while it is being prepared for use. When ready, it is automatically attached to the data center.

APPENDIX A. CHANGING THE PERMISSIONS FOR THE LOCAL ISO DOMAIN

If the Manager was configured during setup to provide a local ISO domain, that domain can be attached to one or more data centers, and used to provide virtual machine image files. By default, the access control list (ACL) for the local ISO domain provides read and write access for only the Manager machine. Virtualization hosts require read and write access to the ISO domain in order to attach the domain to a data center. Use this procedure if network or host details were not available at the time of setup, or if you need to update the ACL at any time.

While it is possible to allow read and write access to the entire network, it is recommended that you limit access to only those hosts and subnets that require it.

Procedure A.1. Changing the Permissions for the Local ISO Domain

1. Log in to the Manager machine.
2. Edit the `/etc/exports` file, and add the hosts, or the subnets to which they belong, to the access control list:

```
/var/lib/exports/iso 10.1.2.0/255.255.255.0(rw)
host01.example.com(rw) host02.example.com(rw)
```

The example above allows read and write access to a single /24 network and two specific hosts. `/var/lib/exports/iso` is the default file path for the ISO domain. See the `exports(5)` man page for further formatting options.

3. Apply the changes:

```
# exportfs -ra
```

Note that if you manually edit the `/etc/exports` file after running `engine-setup`, running `engine-cleanup` later will not undo the changes.

APPENDIX B. ATTACHING THE LOCAL ISO DOMAIN TO A DATA CENTER

The local ISO domain, created during the Manager installation, appears in the Administration Portal as **Unattached**. To use it, attach it to a data center. The ISO domain must be of the same **Storage Type** as the data center. Each host in the data center must have read and write access to the ISO domain. In particular, ensure that the Storage Pool Manager has access.

Only one ISO domain can be attached to a data center.

Procedure B.1. Attaching the Local ISO Domain to a Data Center

1. In the Administration Portal, click the **Data Centers** resource tab and select the appropriate data center.
2. Select the **Storage** tab in the details pane to list the storage domains already attached to the data center.
3. Click **Attach ISO** to open the **Attach ISO Library** window.
4. Click the radio button for the local ISO domain.
5. Click **OK**.

The ISO domain is now attached to the data center and is automatically activated.

APPENDIX C. ENABLING GLUSTER PROCESSES ON RED HAT GLUSTER STORAGE NODES

1. In the Navigation Pane, select the **Clusters** tab.
2. Select **New**.
3. Select the "Enable Gluster Service" radio button. Provide the address, SSH fingerprint, and password as necessary. The address and password fields can be filled in only when the **Import existing Gluster configuration** check box is selected.

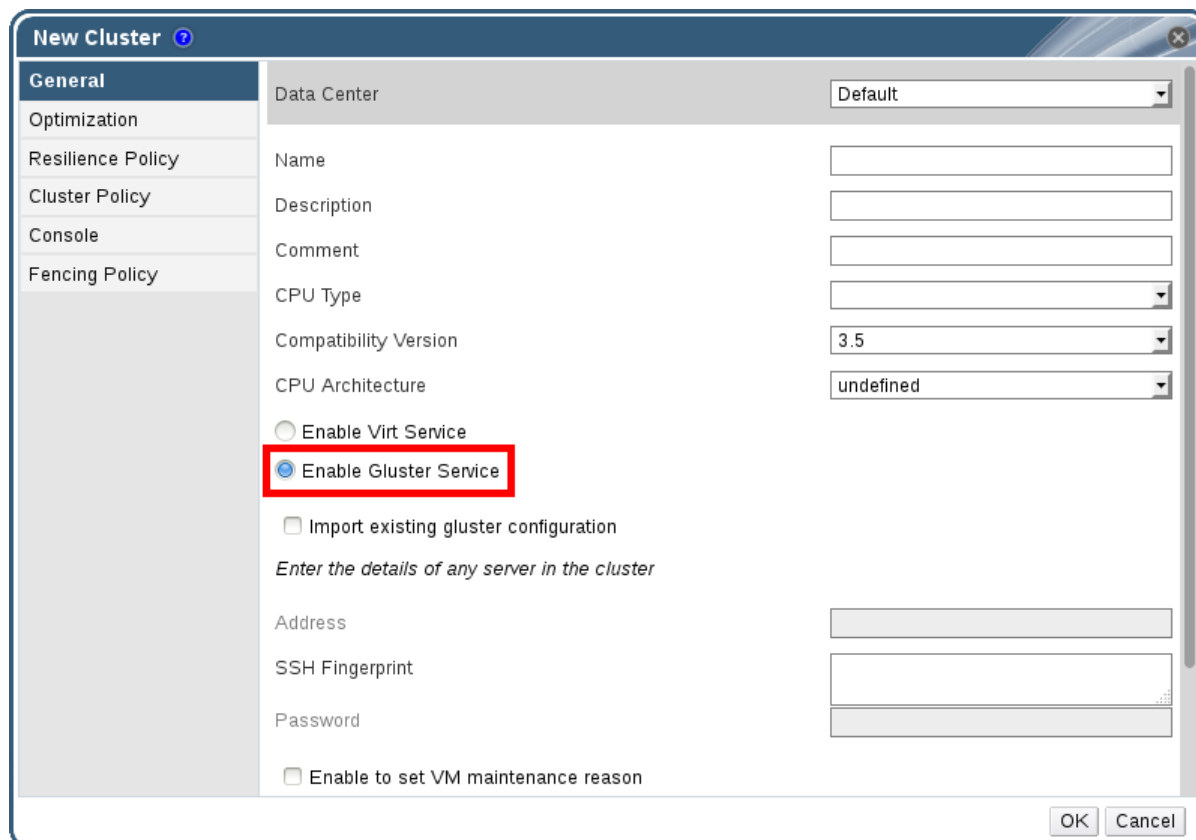


Figure C.1. Selecting the "Enable Gluster Service" Radio Button

4. Click **OK**.

It is now possible to add Red Hat Gluster Storage nodes to the Gluster cluster, and to mount Gluster volumes as storage domains. **iptables** rules no longer block storage domains from being added to the cluster.

APPENDIX D. PREPARING A REMOTE POSTGRESQL DATABASE

You can configure a PostgreSQL database on a remote Red Hat Enterprise Linux 7 machine for the Red Hat Virtualization Manager or for Data Warehouse.

By default, the Manager's configuration script, **engine-setup**, creates and configures the Manager database locally on the Manager machine. For automatic database configuration, see [Section 2.3, “Configuring the Red Hat Virtualization Manager”](#). To set up the Manager database with custom values on the Manager machine, see [Appendix E, *Preparing a Local Manually-Configured PostgreSQL Database for Use with the Red Hat Virtualization Manager*](#). You should set up a Manager database before you configure the Manager; you must supply the database credentials during **engine-setup**.

The Data Warehouse's [configuration script](#) offers the choice of creating a local or remote database. However, situations may arise where you might want to prepare a remote database for Data Warehouse manually.

Use this procedure to configure the database on a machine that is separate from the machine where the Manager is installed.



NOTE

The **engine-setup** and **engine-backup --mode=restore** commands only support system error messages in the **en_US.UTF8** locale, even if the system locale is different.

The locale settings in the **postgresql.conf** file must be set to **en_US.UTF8**.



IMPORTANT

The database name must contain only numbers, underscores, and lowercase letters.

Procedure D.1. Preparing a Remote PostgreSQL Database

1. Install the PostgreSQL server package:

```
# yum install postgresql-server
```

2. Initialize the PostgreSQL database, start the **postgresql** service, and ensure that this service starts on boot:

```
# su -l postgres -c "/usr/bin/initdb --locale=en_US.UTF8 --
auth='ident' --pgdata=/var/lib/pgsql/data/"
# systemctl start postgresql.service
# systemctl enable postgresql.service
```

3. Connect to the **psql** command line interface as the **postgres** user:

```
# su - postgres
$ psql
```

4. Create a default user. The Manager's default user is **engine** and the Data Warehouse's default user is **ovirt_engine_history**:

```
postgres=# create role user_name with login encrypted password
'password' ;
```

5. Create a database. The Manager's default database name is **engine** and Data Warehouse's default database name is **ovirt_engine_history**:

```
postgres=# create database database_name owner user_name template
template0 encoding 'UTF8' lc_collate 'en_US.UTF-8' lc_ctype
'en_US.UTF-8' ;
```

6. Connect to the new database and add the **plpgsql** language:

```
postgres=# \c database_name
database_name=# CREATE LANGUAGE plpgsql;
```

7. Ensure the database can be accessed remotely by enabling md5 client authentication. Edit the **/var/lib/pgsql/data/pg_hba.conf** file, and add the following line immediately underneath the line starting with **local** at the bottom of the file, replacing **X.X.X.X** with the IP address of the Manager or the Data Warehouse machine:

```
host    database_name    user_name    X.X.X.X/32    md5
```

8. Allow TCP/IP connections to the database. Edit the **/var/lib/pgsql/data/postgresql.conf** file and add the following line:

```
listen_addresses='*'
```

This example configures the **postgresql** service to listen for connections on all interfaces. You can specify an interface by giving its IP address.

9. Open the default port used for PostgreSQL database connections, and save the updated firewall rules:

```
# yum install iptables-services
# iptables -I INPUT 5 -p tcp --dport 5432 -j ACCEPT
# service iptables save
```

10. Edit the following parameters in the **/var/lib/pgsql/data/postgresql.conf** file:

```
autovacuum_vacuum_scale_factor='0.01'
autovacuum_analyze_scale_factor='0.075'
autovacuum_max_workers='6'
maintenance_work_mem='65536'
max_connections='150'
```

11. Restart the **postgresql** service:

```
# systemctl restart postgresql.service
```

Optionally, set up SSL to secure database connections using the instructions at <http://www.postgresql.org/docs/9.6/static/ssl-tcp.html#SSL-FILE-USAGE>.

APPENDIX E. PREPARING A LOCAL MANUALLY-CONFIGURED POSTGRESQL DATABASE FOR USE WITH THE RED HAT VIRTUALIZATION MANAGER

Optionally configure a local PostgreSQL database on the Manager machine to use as the Manager database. By default, the Red Hat Virtualization Manager's configuration script, **engine-setup**, creates and configures the Manager database locally on the Manager machine. For automatic database configuration, see [Section 2.3, “Configuring the Red Hat Virtualization Manager”](#). To configure the Manager database on a machine that is separate from the machine where the Manager is installed, see [Appendix D, *Preparing a Remote PostgreSQL Database*](#).

Use this procedure to set up the Manager database with custom values. Set up this database before you configure the Manager; you must supply the database credentials during **engine-setup**. To set up the database, you must first install the `hevm` package on the Manager machine; the `postgresql-server` package is installed as a dependency.



NOTE

The **engine-setup** and **engine-backup --mode=restore** commands only support system error messages in the `en_US.UTF8` locale, even if the system locale is different.

The locale settings in the `postgresql.conf` file must be set to `en_US.UTF8`.



IMPORTANT

The database name must contain only numbers, underscores, and lowercase letters.

Procedure E.1. Preparing a Local Manually-Configured PostgreSQL Database for use with the Red Hat Virtualization Manager

1. Initialize the PostgreSQL database, start the `postgresql` service, and ensure that this service starts on boot:

```
# su -l postgres -c "/usr/bin/initdb --locale=en_US.UTF8 --
auth='ident' --pgdata=/var/lib/pgsql/data/"
# systemctl start postgresql.service
# systemctl enable postgresql.service
```

2. Connect to the `psql` command line interface as the `postgres` user:

```
# su - postgres
$ psql
```

3. Create a user for the Manager to use when it writes to and reads from the database. The default user name on the Manager is **engine**:

```
postgres=# create role user_name with login encrypted password
'password';
```

4. Create a database in which to store data about the Red Hat Virtualization environment. The default database name on the Manager is **engine**:

```
postgres=# create database database_name owner user_name template
template0 encoding 'UTF8' lc_collate 'en_US.UTF-8' lc_ctype
'en_US.UTF-8';
```

5. Connect to the new database and add the **plpgsql** language:

```
postgres=# \c database_name
database_name=# CREATE LANGUAGE plpgsql;
```

6. Ensure the database can be accessed remotely by enabling md5 client authentication. Edit the `/var/lib/pgsql/data/pg_hba.conf` file, and add the following line immediately underneath the line starting with **local** at the bottom of the file:

```
host      [database name]      [user name]      0.0.0.0/0      md5
host      [database name]      [user name]      :::0/0         md5
```

7. Restart the **postgresql** service:

```
# systemctl restart postgresql.service
```

Optionally, set up SSL to secure database connections using the instructions at <http://www.postgresql.org/docs/8.4/static/ssl-tcp.html#SSL-FILE-USAGE>.

APPENDIX F. INSTALLING A WEBSOCKET PROXY ON A SEPARATE MACHINE

The websocket proxy allows users to connect to virtual machines via noVNC and SPICE HTML5 consoles. The noVNC client uses websockets to pass VNC data. However, the VNC server in QEMU does not provide websocket support, therefore a websocket proxy must be placed between the client and the VNC server. The proxy can run on any machine that has access to the network, including the the Manager machine.

For security and performance reasons, users may want to configure the websocket proxy on a separate machine.



NOTE

SPICE HTML5 support is a Technology Preview feature. Technology Preview features are not fully supported under Red Hat Subscription Service Level Agreements (SLAs), may not be functionally complete, and are not intended for production use. However, these features provide early access to upcoming product innovations, enabling customers to test functionality and provide feedback during the development process.

This section describes how to install and configure the websocket proxy on a separate machine that does not run the Manager. See [Section 2.3, “Configuring the Red Hat Virtualization Manager”](#) for instructions on how to configure the websocket proxy on the Manager.


Procedure F.1. Installing and Configuring a Websocket Proxy on a Separate Machine

1. Install the websocket proxy:

```
# yum install ovirt-engine-websocket-proxy
```

2. Run the **engine-setup** command to configure the websocket proxy.

```
# engine-setup
```



NOTE

If the `rhev` package has also been installed, choose **No** when asked to configure the engine on this host.

3. Press **Enter** to allow **engine-setup** to configure a websocket proxy server on the machine.

```
Configure WebSocket Proxy on this machine? (Yes, No) [Yes]:
```

4. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**. Note that the automatically detected hostname may be incorrect if you are using virtual hosts:

Host fully qualified DNS name of this server [*host.example.com*]:

5. Press **Enter** to allow **engine-setup** to configure the firewall and open the ports required for external communication. If you do not allow **engine-setup** to modify your firewall configuration, then you must manually open the required ports.

Setup can automatically configure the firewall on this system.
Note: automatic configuration of the firewall may overwrite current settings.

Do you want Setup to configure the firewall? (Yes, No) [Yes]:

6. Enter the fully qualified DNS name of the Manager machine and press **Enter**.

Host fully qualified DNS name of the engine server []:
engine_host.example.com

7. Press **Enter** to allow **engine-setup** to perform actions on the Manager machine, or press **2** to manually perform the actions.

Setup will need to do some actions on the remote engine server.
Either automatically, using ssh as root to access it, or you will be prompted to manually perform each such action.

Please choose one of the following:

- 1 - Access remote engine server using ssh as root
 - 2 - Perform each action manually, use files to copy content around
- (1, 2) [1]:

- a. Press **Enter** to accept the default SSH port number, or enter the port number of the Manager machine.

ssh port on remote engine server [22]:

- b. Enter the root password to log in to the Manager machine and press **Enter**.

root password on remote engine server *engine_host.example.com*:

8. Select whether to review iptables rules if they differ from the current settings.

Generated iptables rules are different from current ones.
Do you want to review them? (Yes, No) [No]:

9. Press **Enter** to confirm the configuration settings.

--== CONFIGURATION PREVIEW ==--

```

Firewall manager           : iptables
Update Firewall           : True
Host FQDN                  : host.example.com
Configure WebSocket Proxy  : True
Engine Host FQDN           : engine_host.example.com

```

Please confirm installation settings (OK, Cancel) [OK]:

Instructions are provided to configure the Manager machine to use the configured websocket proxy.

```
Manual actions are required on the engine host
in order to enroll certs for this host and configure the engine
about it.
```

```
Please execute this command on the engine host:
  engine-config -s WebSocketProxy=host.example.com:6100
and than restart the engine service to make it effective
```

10. Log in to the Manager machine and execute the provided instructions.

```
# engine-config -s WebSocketProxy=host.example.com:6100
# systemctl restart ovirt-engine.service
```

APPENDIX G. CONFIGURING A HOST FOR PCI PASSTHROUGH

Enabling PCI passthrough allows a virtual machine to use a host device as if the device were directly attached to the virtual machine. To enable the PCI passthrough function, you need to enable virtualization extensions and the IOMMU function. The following procedure requires you to reboot the host. If the host is attached to the Manager already, ensure you place the host into maintenance mode before running the following procedure.

Prerequisites:

- Ensure that the host hardware meets the requirements for PCI device passthrough and assignment. See [PCI Device Requirements](#) in the *Planning and Prerequisites Guide* for more information.

Procedure G.1. Configuring a Host for PCI Passthrough

1. Enable the virtualization extension and IOMMU extension in the BIOS. See [Enabling Intel VT-x and AMD-V virtualization hardware extensions in BIOS](#) in the *Red Hat Enterprise Linux Virtualization and Administration Guide* for more information.
2. Enable the IOMMU flag in the kernel by selecting the **Hostdev Passthrough & SR-IOV** check box when adding the host to the Manager or by editing the `grub` configuration file manually.
 - To enable the IOMMU flag from the Administration Portal, see [Adding a Host to the Red Hat Virtualization Manager](#) and [Kernel Settings Explained](#) in the *Administration Guide*.
 - To edit the `grub` configuration file manually, see [Procedure G.2, “Enabling IOMMU Manually”](#).
3. For GPU passthrough, you need to run additional configuration steps on both the host and the guest system. See [Preparing Host and Guest Systems for GPU Passthrough](#) in the *Administration Guide* for more information.

Procedure G.2. Enabling IOMMU Manually

1. Enable IOMMU by editing the `grub` configuration file.



NOTE

If you are using IBM POWER8 hardware, skip this step as IOMMU is enabled by default.

- For Intel, boot the machine, and append `intel_iommu=on` to the end of the `GRUB_CMDLINE_LINUX` line in the `grub` configuration file.

```
# vi /etc/default/grub
...
GRUB_CMDLINE_LINUX="nofb splash=quiet console=tty0 ..."
intel_iommu=on
...
```

- For AMD, boot the machine, and append **amd_iommu=on** to the end of the **GRUB_CMDLINE_LINUX** line in the **grub** configuration file.

```
# vi /etc/default/grub
...
GRUB_CMDLINE_LINUX="nofb splash=quiet console=tty0 ..."
amd_iommu=on
...
```

NOTE

If **intel_iommu=on** or **amd_iommu=on** works, you can try replacing them with **iommu=pt** or **amd_iommu=pt**. The **pt** option only enables IOMMU for devices used in passthrough and will provide better host performance. However, the option may not be supported on all hardware. Revert to previous option if the **pt** option doesn't work for your host.

If the passthrough fails because the hardware does not support interrupt remapping, you can consider enabling the **allow_unsafe_interrupts** option if the virtual machines are trusted. The **allow_unsafe_interrupts** is not enabled by default because enabling it potentially exposes the host to MSI attacks from virtual machines. To enable the option:

```
# vi /etc/modprobe.d
options vfio_iommu_type1 allow_unsafe_interrupts=1
```

2. Refresh the **grub.cfg** file and reboot the host for these changes to take effect:

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

```
# reboot
```

For enabling SR-IOV and assigning dedicated virtual NICs to virtual machines, see <https://access.redhat.com/articles/2335291> for more information.