



Red Hat Virtualization 4.3

Migrating from a standalone Manager to a self-hosted engine

How to migrate the Red Hat Virtualization Manager from a standalone server to a self-managed virtual machine

Red Hat Virtualization 4.3 Migrating from a standalone Manager to a self-hosted engine

How to migrate the Red Hat Virtualization Manager from a standalone server to a self-managed virtual machine

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

Legal Notice

Copyright © 2022 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 how to migrate your Red Hat Virtualization environment from a Manager installed on a physical server (or a virtual machine running in a separate environment) to a Manager installed on a virtual machine running on the same hosts it manages.

Table of Contents

PREFACE	3
Standalone Manager Architecture	3
Self-Hosted Engine Architecture	3
CHAPTER 1. MIGRATION OVERVIEW	5
CHAPTER 2. INSTALLING THE SELF-HOSTED ENGINE DEPLOYMENT HOST	6
2.1. INSTALLING RED HAT VIRTUALIZATION HOSTS	6
2.1.1. Enabling the Red Hat Virtualization Host Repository	7
2.2. INSTALLING RED HAT ENTERPRISE LINUX HOSTS	8
2.2.1. Enabling the Red Hat Enterprise Linux host Repositories	9
CHAPTER 3. PREPARING STORAGE FOR RED HAT VIRTUALIZATION	11
3.1. PREPARING NFS STORAGE	11
3.2. PREPARING ISCSI STORAGE	12
3.3. PREPARING FCP STORAGE	13
3.4. PREPARING RED HAT GLUSTER STORAGE	14
3.5. CUSTOMIZING MULTIPATH CONFIGURATIONS FOR SAN VENDORS	14
3.6. RECOMMENDED SETTINGS FOR MULTIPATH.CONF	15
CHAPTER 4. UPDATING THE RED HAT VIRTUALIZATION MANAGER	17
CHAPTER 5. BACKING UP THE ORIGINAL MANAGER	18
CHAPTER 6. RESTORING THE BACKUP ON A NEW SELF-HOSTED ENGINE	19
CHAPTER 7. ENABLING THE RED HAT VIRTUALIZATION MANAGER REPOSITORIES	22
CHAPTER 8. REINSTALLING AN EXISTING HOST AS A SELF-HOSTED ENGINE NODE	24

PREFACE

You can convert a standalone Red Hat Virtualization Manager to a self-hosted engine by backing up the standalone Manager and restoring it in a new self-hosted environment.

The difference between the two environment types is explained below:


Standalone Manager Architecture

The Red Hat Virtualization Manager runs on a physical server, or a virtual machine hosted in a separate virtualization environment. A standalone Manager is easier to deploy and manage, but requires an additional physical server. The Manager is only highly available when managed externally with a product such as Red Hat's High Availability Add-On.

The minimum setup for a standalone Manager environment includes:

- One Red Hat Virtualization Manager machine. The Manager is typically deployed on a physical server. However, it can also be deployed on a virtual machine, as long as that virtual machine is hosted in a separate environment. The Manager must run on Red Hat Enterprise Linux 7.
- A minimum of two hosts for virtual machine high availability. You can use Red Hat Enterprise Linux hosts or Red Hat Virtualization Hosts (RHVH). VDSM (the host agent) runs on all hosts to facilitate communication with the Red Hat Virtualization Manager.
- One storage service, which can be hosted locally or on a remote server, depending on the storage type used. The storage service must be accessible to all hosts.

Figure 1. Standalone Manager Red Hat Virtualization Architecture




Self-Hosted Engine Architecture

The Red Hat Virtualization Manager runs as a virtual machine on self-hosted engine nodes (specialized hosts) in the same environment it manages. A self-hosted engine environment requires one less physical server, but requires more administrative overhead to deploy and manage. The Manager is highly available without external HA management.

The minimum setup of a self-hosted engine environment includes:

- One Red Hat Virtualization Manager virtual machine that is hosted on the self-hosted engine nodes. The RHV-M Appliance is used to automate the installation of a Red Hat Enterprise Linux 7 virtual machine, and the Manager on that virtual machine.
- A minimum of two self-hosted engine nodes for virtual machine high availability. You can use Red Hat Enterprise Linux hosts or Red Hat Virtualization Hosts (RHVH). VDSM (the host agent) runs on all hosts to facilitate communication with the Red Hat Virtualization Manager. The HA services run on all self-hosted engine nodes to manage the high availability of the Manager virtual machine.
- One storage service, which can be hosted locally or on a remote server, depending on the storage type used. The storage service must be accessible to all hosts.

Figure 2. Self-Hosted Engine Red Hat Virtualization Architecture



CHAPTER 1. MIGRATION OVERVIEW

When you specify a backup file during self-hosted engine deployment, the Manager backup is restored on a new virtual machine, with a dedicated self-hosted engine storage domain. Deploying on a fresh host is highly recommended; if the host used for deployment existed in the backed up environment, it will be removed from the restored database to avoid conflicts in the new environment.

At least two self-hosted engine nodes are required for the Manager virtual machine to be highly available. You can add new nodes, or convert existing hosts.

The migration involves the following key steps:

1. [Install a new host to deploy the self-hosted engine on.](#) You can use either host type:
 - [Red Hat Virtualization Host](#)
 - [Red Hat Enterprise Linux](#)
2. [Prepare storage for the self-hosted engine storage domain.](#) You can use one of the following storage types:
 - [NFS](#)
 - [iSCSI](#)
 - [Fibre Channel \(FCP\)](#)
 - [Red Hat Gluster Storage](#)
3. [Update the original Manager to the latest minor version before you back it up.](#)
4. [Back up the original Manager using the **engine-backup** tool.](#)
5. [Deploy a new self-hosted engine and restore the backup.](#)
6. [Enable the Manager repositories on the new Manager virtual machine.](#)
7. [Convert regular hosts to self-hosted engine nodes that can host the new Manager.](#)


This procedure assumes that you have access and can make changes to the original Manager.

Prerequisites

- FQDNs prepared for your Manager and the deployment host. Forward and reverse lookup records must both be set in the DNS. The new Manager must have the same FQDN as the original Manager.
- The management network (**ovirtmgmt** by default) must be configured as a **VM network**, so that it can manage the Manager virtual machine.

CHAPTER 2. INSTALLING THE SELF-HOSTED ENGINE DEPLOYMENT HOST

A self-hosted engine can be deployed from a [Red Hat Virtualization Host](#) or a [Red Hat Enterprise Linux host](#).



IMPORTANT

If you plan to use bonded interfaces for high availability or VLANs to separate different types of traffic (for example, for storage or management connections), you should configure them on the host before beginning the self-hosted engine deployment. See [Networking Recommendations](#) in the *Planning and Prerequisites Guide*.

2.1. INSTALLING RED HAT VIRTUALIZATION HOSTS

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 web interface for monitoring the host and performing administrative tasks. See <http://cockpit-project.org/running.html> for the minimum browser requirements.

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

The host must meet the minimum [host requirements](#).

Procedure


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. Go to **Hypervisor Image for RHV 4.3** 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 you are installing RHVH, booting from the prepared installation media.
3. From the boot menu, select **Install RHVH 4.3** 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.

8. Select a network from the **Network & Host Name** screen and click **Configure...** to configure the connection details.



NOTE

To use the connection every time the system boots, select the **Automatically connect to this network when it is available** check box. For more information, see [Edit Network Connections](#) in the *Red Hat Enterprise Linux 7 Installation Guide*.

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.

2.1.1. Enabling the Red Hat Virtualization Host Repository

Register the system to receive updates. Red Hat Virtualization Host only requires one repository. This section provides instructions for registering RHVH with the [Content Delivery Network](#), or with [Red Hat Satellite 6](#).

Registering RHVH with the Content Delivery Network

1. Log in to the Cockpit web interface at **`https://HostFQDNorIP:9090`**.
2. Navigate to **Subscriptions**, click **Register System**, and enter your Customer Portal user name and password. The **Red Hat Virtualization Host** subscription is automatically attached to the system.
3. Click **Terminal**.
4. 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
```

Registering RHVH with Red Hat Satellite 6

1. Log in to the Cockpit web interface at **`https://HostFQDNorIP:9090`**.
2. Click **Terminal**.
3. Register RHVH with Red Hat Satellite 6:

```
# rpm -Uvh http://satellite.example.com/pub/katello-ca-consumer-latest.noarch.rpm
# subscription-manager register --org="org_id"
# subscription-manager list --available
# subscription-manager attach --pool=pool_id
# subscription-manager repos \
--disable='*' \
--enable=rhel-7-server-rhvh-4-rpms
```

2.2. INSTALLING RED HAT ENTERPRISE LINUX HOSTS

A Red Hat Enterprise Linux host is based on a standard basic installation of Red Hat Enterprise Linux 7 on a physical server, with the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** subscriptions attached.

For detailed installation instructions, see the [Performing a standard {enterprise-linux-shortname} installation](#).

The host must meet the minimum [host requirements](#).



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.

2.2.1. Enabling the Red Hat Enterprise Linux host Repositories

To use a Red Hat Enterprise Linux machine as a host, you must register the system with the Content Delivery Network, attach the **Red Hat Enterprise Linux Server** and **Red Hat Virtualization** subscriptions, and enable the host repositories.

Procedure

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 record the pool IDs:

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

3. Use the pool IDs to attach the subscriptions to the system:

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



NOTE

To view currently attached subscriptions:

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

To list all enabled repositories:

```
# yum repolist
```

4. Configure the repositories:

```
# subscription-manager repos \
  --disable='*' \
  --enable=rhel-7-server-rpms \
  --enable=rhel-7-server-rhv-4-mgmt-agent-rpms \
  --enable=rhel-7-server-ansible-2.9-rpms
```

For Red Hat Enterprise Linux 7 hosts, little endian, on IBM POWER8 hardware:

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

For Red Hat Enterprise Linux 7 hosts, little endian, on IBM POWER9 hardware:

```
# subscription-manager repos \  
--disable='*' \  
--enable=rhel-7-server-rhv-4-mgmt-agent-for-power-9-rpms \  
--enable=rhel-7-for-power-9-rpms
```

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

```
# yum update
```

6. Reboot the machine.

Although the existing storage domains will be migrated from the standalone Manager, you must prepare additional storage for a self-hosted engine storage domain that is dedicated to the Manager virtual machine.

CHAPTER 3. PREPARING STORAGE FOR RED HAT VIRTUALIZATION

Prepare storage to be used for storage domains in the new environment. A Red Hat Virtualization environment must have at least one data storage domain, but adding more is recommended.

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 while active (but can be migrated between data centers). Data domains of multiple storage types can be added to the same data center, provided they are all shared, rather than local, domains.

Self-hosted engines must have an additional data domain dedicated to the Manager virtual machine. This domain is created during the self-hosted engine deployment, and must be at least 74 GiB. You must prepare the storage for this domain before beginning the deployment.

You can use one of the following storage types:

- [NFS](#)
- [iSCSI](#)
- [Fibre Channel \(FCP\)](#)
- [Red Hat Gluster Storage](#)



IMPORTANT

If you are using iSCSI storage, the self-hosted engine storage domain must use its own iSCSI target. Any additional storage domains must use a different iSCSI target.



WARNING

Creating additional data storage domains in the same data center as the self-hosted engine storage domain is highly recommended. If you deploy the self-hosted engine in a data center with only one active data storage domain, and that storage domain is corrupted, you will not be able to add new storage domains or remove the corrupted storage domain; you will have to redeploy the self-hosted engine.

3.1. PREPARING NFS STORAGE

Set up NFS shares on your file storage or remote server to serve as storage domains on Red Hat Enterprise Virtualization Host systems. After exporting the shares on the remote storage and configuring them in the Red Hat Virtualization Manager, the shares will be automatically imported on the Red Hat Virtualization hosts.

For information on setting up and configuring NFS, see [Network File System \(NFS\)](#) in the *Red Hat Enterprise Linux 7 Storage Administration Guide*.

For information on how to export an 'NFS' share, see [How to export 'NFS' share from NetApp Storage / EMC SAN in Red Hat Virtualization](#)

Specific system user accounts and system user groups are required by Red Hat Virtualization so the Manager can store data in the storage domains represented by the exported directories. The following procedure sets the permissions for one directory. You must repeat the **chown** and **chmod** steps for all of the directories you intend to use as storage domains in Red Hat Virtualization.

Procedure

1. Create the group **kvm**:

```
# groupadd kvm -g 36
```

2. Create the user **vds**m in the group **kvm**:

```
# useradd vds -u 36 -g 36
```

3. Set the ownership of your exported directory to 36:36, which gives **vds**m:**kvm** ownership:

```
# chown -R 36:36 /exports/data
```

4. Change the mode of the directory so that read and write access is granted to the owner, and so that read and execute access is granted to the group and other users:

```
# chmod 0755 /exports/data
```

3.2. PREPARING ISCSI STORAGE


Red Hat Virtualization supports iSCSI storage, which is a storage domain created from a volume group made up of LUNs. Volume groups and LUNs cannot be attached to more than one storage domain at a time.

For information on setting up and configuring iSCSI storage, see [Online Storage Management](#) in the *Red Hat Enterprise Linux 7 Storage Administration Guide* .



IMPORTANT

If you are using block storage and you intend to deploy virtual machines on raw devices or direct LUNs and to manage them with the Logical Volume Manager, you must create a filter to hide the guest logical volumes. This will prevent guest logical volumes from being activated when the host is booted, a situation that could lead to stale logical volumes and cause data corruption. See <https://access.redhat.com/solutions/2662261> for details.



IMPORTANT

Red Hat Virtualization currently does not support block storage with a block size of 4K. You must configure block storage in legacy (512b block) mode.

IMPORTANT

If your host is booting from SAN storage and loses connectivity to the storage, the storage file systems become read-only and remain in this state after connectivity is restored.

To prevent this situation, Red Hat recommends adding a drop-in multipath configuration file on the root file system of the SAN for the boot LUN to ensure that it is queued when there is a connection:

```
# cat /etc/multipath/conf.d/host.conf
multipaths {
  multipath {
    wwid boot_LUN_wwid
    no_path_retry queue
  }
}
```

3.3. PREPARING FCP STORAGE

Red Hat Virtualization 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 on setting up and configuring FCP or multipathing on Red Hat Enterprise Linux, see the [Storage Administration Guide](#) and [DM Multipath Guide](#).

IMPORTANT

If you are using block storage and you intend to deploy virtual machines on raw devices or direct LUNs and to manage them with the Logical Volume Manager, you must create a filter to hide the guest logical volumes. This will prevent guest logical volumes from being activated when the host is booted, a situation that could lead to stale logical volumes and cause data corruption. See <https://access.redhat.com/solutions/2662261> for details.

IMPORTANT

Red Hat Virtualization currently does not support block storage with a block size of 4K. You must configure block storage in legacy (512b block) mode.



IMPORTANT

If your host is booting from SAN storage and loses connectivity to the storage, the storage file systems become read-only and remain in this state after connectivity is restored.

To prevent this situation, Red Hat recommends adding a drop-in multipath configuration file on the root file system of the SAN for the boot LUN to ensure that it is queued when there is a connection:

```
# cat /etc/multipath/conf.d/host.conf
multipaths {
  multipath {
    wwid boot_LUN_wwid
    no_path_retry queue
  }
}
```

3.4. PREPARING RED HAT GLUSTER STORAGE

For information on setting up and configuring Red Hat Gluster Storage, see the [Red Hat Gluster Storage Installation Guide](#).

For the Red Hat Gluster Storage versions that are supported with Red Hat Virtualization, see <https://access.redhat.com/articles/2356261>.

3.5. CUSTOMIZING MULTIPATH CONFIGURATIONS FOR SAN VENDORS

To customize the multipath configuration settings, do not modify **/etc/multipath.conf**. Instead, create a new configuration file that overrides **/etc/multipath.conf**.



WARNING

Upgrading Virtual Desktop and Server Manager (VDSM) overwrites the **/etc/multipath.conf** file. If **multipath.conf** contains customizations, overwriting it can trigger storage issues.

Prerequisites

- This topic only applies to systems that have been configured to use multipath connections storage domains, and therefore have a **/etc/multipath.conf** file.
- Do not override the **user_friendly_names** and **find_multipaths** settings. For more information, see [Section 3.6, "Recommended Settings for Multipath.conf"](#)
- Avoid overriding **no_path_retry** and **polling_interval** unless required by the storage vendor. For more information, see [Section 3.6, "Recommended Settings for Multipath.conf"](#)

Procedure

1. To override the values of settings in **/etc/multipath.conf**, create a new configuration file in the **/etc/multipath/conf.d/** directory.



NOTE

The files in **/etc/multipath/conf.d/** execute in alphabetical order. Follow the convention of naming the file with a number at the beginning of its name. For example, **/etc/multipath/conf.d/90-myfile.conf**.

2. Copy the settings you want to override from **/etc/multipath.conf** to the new configuration file in **/etc/multipath/conf.d/**. Edit the setting values and save your changes.
3. Apply the new configuration settings by entering the **systemctl reload multipathd** command.



NOTE

Avoid restarting the multipathd service. Doing so generates errors in the VDSM logs.

Verification steps

If you override the VDSM-generated settings in **/etc/multipath.conf**, verify that the new configuration performs as expected in a variety of failure scenarios.

For example, disable all of the storage connections. Then enable one connection at a time and verify that doing so makes the storage domain reachable.

Troubleshooting

If a Red Hat Virtualization Host has trouble accessing shared storage, check **/etc/multipath.conf** and files under **/etc/multipath/conf.d/** for values that are incompatible with the SAN.

Additional resources

- [Red Hat Enterprise Linux DM Multipath](#) in the RHEL documentation.
- [Configuring iSCSI Multipathing](#) in the Administration Guide.
- [How do I customize /etc/multipath.conf on my RHVH hypervisors? What values must not change and why?](#) on the Red Hat Customer Portal, which shows an example **multipath.conf** file and was the basis for this topic.

3.6. RECOMMENDED SETTINGS FOR MULTIPATH.CONF

When overriding **/etc/multipath.conf**, Do not override the following settings:

user_friendly_names no

This setting controls whether user-friendly names are assigned to devices in addition to the actual device names. Multiple hosts must use the same name to access devices. Disabling this setting prevents user-friendly names from interfering with this requirement.

find_multipaths no

This setting controls whether RHVH tries to access all devices through multipath, even if only one path is available. Disabling this setting prevents RHV from using the too-clever behavior when this setting is enabled.

Avoid overriding the following settings unless required by the storage system vendor:

no_path_retry 4

This setting controls the number of polling attempts to retry when no paths are available. Before RHV version 4.2, the value of **no_path_retry** was **fail** because QEMU had trouble with the I/O queuing when no paths were available. The **fail** value made it fail quickly and paused the virtual machine. RHV version 4.2 changed this value to **4** so when multipathd detects the last path has failed, it checks all of the paths four more times. Assuming the default 5-second polling interval, checking the paths takes 20 seconds. If no path is up, multipathd tells the kernel to stop queuing and fails all outstanding and future I/O until a path is restored. When a path is restored, the 20-second delay is reset for the next time all paths fail. For more details, see [the commit that changed this setting](#).

polling_interval 5

This setting determines the number of seconds between polling attempts to detect whether a path is open or has failed. Unless the vendor provides a clear reason for increasing the value, keep the VDSM-generated default so the system responds to path failures sooner.

Before backing up the Manager, ensure it is updated to the latest minor version. The Manager version in the backup file must match the version of the new Manager.

CHAPTER 4. UPDATING THE RED HAT VIRTUALIZATION MANAGER

Updates to the Red Hat Virtualization Manager are released through the Content Delivery Network.

Procedure

1. Check if updated packages are available:

```
# engine-upgrade-check
```

2. Update the setup packages:

```
# yum update ovirt\*setup\* rh\*vm-setup-plugins
```

3. Update the Red Hat Virtualization Manager with the **engine-setup** script. The **engine-setup** script prompts you with some configuration questions, then stops the **ovirt-engine** service, downloads and installs the updated packages, backs up and updates the database, performs post-installation configuration, and starts the **ovirt-engine** service.

```
# engine-setup
```

When the script completes successfully, the following message appears:

```
Execution of setup completed successfully
```



NOTE

The **engine-setup** script is also used during the Red Hat Virtualization Manager installation process, and it stores the configuration values supplied. During an update, the stored values are displayed when previewing the configuration, and might not be up to date if **engine-config** was used to update configuration after installation. For example, if **engine-config** was used to update **SANWipeAfterDelete** to **true** after installation, **engine-setup** will output "Default SAN wipe after delete: False" in the configuration preview. However, the updated values will not be overwritten by **engine-setup**.



IMPORTANT

The update process might take some time. Do not stop the process before it completes.

4. Update the base operating system and any optional packages installed on the Manager:

```
# yum update
```



IMPORTANT

If any kernel packages were updated, reboot the machine to complete the update.

CHAPTER 5. BACKING UP THE ORIGINAL MANAGER

Back up the original Manager using the **engine-backup** command, and copy the backup file to a separate location so that it can be accessed at any point during the process.

For more information about **engine-backup --mode=backup** options, see [Backing Up and Restoring the Red Hat Virtualization Manager](#) in the *Administration Guide*.

Procedure

1. Log in to the original Manager and stop the **ovirt-engine** service:

```
# systemctl stop ovirt-engine
# systemctl disable ovirt-engine
```



NOTE

Though stopping the original Manager from running is not obligatory, it is recommended as it ensures no changes are made to the environment after the backup is created. Additionally, it prevents the original Manager and the new Manager from simultaneously managing existing resources.

2. Run the **engine-backup** command, specifying the name of the backup file to create, and the name of the log file to create to store the backup log:

```
# engine-backup --mode=backup --file=file_name --log=log_file_name
```

3. Copy the files to an external server. In the following example, **storage.example.com** is the fully qualified domain name of a network storage server that will store the backup until it is needed, and **/backup/** is any designated folder or path.

```
# scp -p file_name log_file_name storage.example.com:/backup/
```


4. If you do not require the Manager machine for other purposes, unregister it from Red Hat Subscription Manager:

```
# subscription-manager unregister
```

After backing up the Manager, deploy a new self-hosted engine and restore the backup on the new virtual machine.

CHAPTER 6. RESTORING THE BACKUP ON A NEW SELF-HOSTED ENGINE

Run the **hosted-engine** script on a new host, and use the **--restore-from-file=path/to/file_name** option to restore the Manager backup during the deployment.



IMPORTANT

If you are using iSCSI storage, and your iSCSI target filters connections according to the initiator's ACL, the deployment may fail with a **STORAGE_DOMAIN_UNREACHABLE** error. To prevent this, you must update your iSCSI configuration before beginning the self-hosted engine deployment:

- If you are redeploying on an existing host, you must update the host's iSCSI initiator settings in **/etc/iscsi/initiatorname.iscsi**. The initiator IQN must be the same as was previously mapped on the iSCSI target, or updated to a new IQN, if applicable.
- If you are deploying on a fresh host, you must update the iSCSI target configuration to accept connections from that host.

Note that the IQN can be updated on the host side (iSCSI initiator), or on the storage side (iSCSI target).

Procedure

1. Copy the backup file to the new host. In the following example, **host.example.com** is the FQDN for the host, and **/backup/** is any designated folder or path.

```
# scp -p file_name host.example.com:/backup/
```

2. Log in to the new host. If you are deploying on Red Hat Virtualization Host, the self-hosted engine deployment tool is available by default. If you are deploying on Red Hat Enterprise Linux, you must install the package:

```
# yum install ovirt-hosted-engine-setup
```

3. Red Hat recommends using the **screen** window manager to run the script to avoid losing the session in case of network or terminal disruption. Install and run **screen**:

```
# yum install screen
# screen
```

In the event of session timeout or connection disruption, run **screen -d -r** to recover the deployment session.

4. Run the **hosted-engine** script, specifying the path to the backup file:

```
# hosted-engine --deploy --restore-from-file=backup/file_name
```

To escape the script at any time, use **CTRL+D** to abort deployment.

5. Select **Yes** to begin the deployment.


6. Configure the network. The script detects possible NICs to use as a management bridge for the environment.
7. If you want to use a custom appliance for the virtual machine installation, enter the path to the OVA archive. Otherwise, leave this field empty to use the RHV-M Appliance.
8. Specify the FQDN for the Manager virtual machine.
9. Enter the root password for the Manager.
10. Enter an SSH public key that will allow you to log in to the Manager as the root user, and specify whether to enable SSH access for the root user.
11. Enter the virtual machine's CPU and memory configuration.



NOTE

The virtual machine must have the same amount of RAM as the physical machine from which the Manager is being migrated. If you must migrate to a virtual machine that has less RAM than the physical machine from which the Manager is migrated, see <https://access.redhat.com/articles/2705841>.

12. Enter a MAC address for the Manager virtual machine, or accept a randomly generated one. If you want to provide the Manager virtual machine with an IP address via DHCP, ensure that you have a valid DHCP reservation for this MAC address. The deployment script will not configure the DHCP server for you.
13. Enter the virtual machine's networking details. If you specify **Static**, enter the IP address of the Manager.



IMPORTANT

The static IP address must belong to the same subnet as the host. For example, if the host is in 10.1.1.0/24, the Manager virtual machine's IP must be in the same subnet range (10.1.1.1-254/24).

14. Specify whether to add entries for the Manager virtual machine and the base host to the virtual machine's **/etc/hosts** file. You must ensure that the host names are resolvable.
15. Provide the name and TCP port number of the SMTP server, the email address used to send email notifications, and a comma-separated list of email addresses to receive these notifications:
16. Enter a password for the **admin@internal** user to access the Administration Portal. The script creates the virtual machine. This can take some time if the RHV-M Appliance needs to be installed.
17. Select the type of storage to use:
 - For NFS, enter the version, full address and path to the storage, and any mount options.
 - For iSCSI, enter the portal details and select a target and LUN from the auto-detected lists. You can only select one iSCSI target during the deployment, but multipathing is supported to connect all portals of the same portal group.



NOTE

To specify more than one iSCSI target, you must enable multipathing before deploying the self-hosted engine. See [Red Hat Enterprise Linux DM Multipath](#) for details. There is also a [Multipath Helper](#) tool that generates a script to install and configure multipath with different options.

- For Gluster storage, enter the full address and path to the storage, and any mount options.



IMPORTANT

Only replica 3 Gluster storage is supported. Ensure you have the following configuration:

- In the `/etc/glusterfs/glusterd.vol` file on all three Gluster servers, set **rpc-auth-allow-insecure** to **on**.

```
option rpc-auth-allow-insecure on
```

- Configure the volume as follows:

```
gluster volume set _volume_ cluster.quorum-type auto
gluster volume set _volume_ network.ping-timeout 10
gluster volume set _volume_ auth.allow \*
gluster volume set _volume_ group virt
gluster volume set _volume_ storage.owner-uid 36
gluster volume set _volume_ storage.owner-gid 36
gluster volume set _volume_ server.allow-insecure on
```

- For Fibre Channel, select a LUN from the auto-detected list. The host bus adapters must be configured and connected, and the LUN must not contain any existing data. To reuse an existing LUN, see [Reusing LUNs](#) in the *Administration Guide*.

18. Enter the Manager disk size.

The script continues until the deployment is complete.

19. The deployment process changes the Manager's SSH keys. To allow client machines to access the new Manager without SSH errors, remove the original Manager's entry from the **.ssh/known_hosts** file on any client machines that accessed the original Manager.

When the deployment is complete, log in to the new Manager virtual machine and enable the required repositories.

CHAPTER 7. ENABLING THE RED HAT VIRTUALIZATION MANAGER REPOSITORIES

Register the system with Red Hat Subscription Manager, attach the **Red Hat Virtualization Manager** subscription, and enable Manager repositories.

Procedure

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

```
# subscription-manager register
```



NOTE

If you are using an IPv6 network, use an IPv6 transition mechanism to access the Content Delivery Network and subscription manager.

2. Find the **Red Hat Virtualization Manager** subscription pool and record the pool ID:

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

3. Use the pool ID to attach the subscription to the system:

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



NOTE

To view currently attached subscriptions:

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

To list all enabled repositories:

```
# yum repolist
```

4. Configure the repositories:

```
# subscription-manager repos \
  --disable='*' \
  --enable=rhel-7-server-rpms \
  --enable=rhel-7-server-supplementary-rpms \
  --enable=rhel-7-server-rhv-4.3-manager-rpms \
  --enable=rhel-7-server-rhv-4-manager-tools-rpms \
  --enable=rhel-7-server-ansible-2.9-rpms \
  --enable=jb-eap-7.2-for-rhel-7-server-rpms
```

The Red Hat Virtualization Manager has been migrated to a self-hosted engine setup. The Manager is now operating on a virtual machine on the new self-hosted engine node.

The hosts will be running in the new environment, but cannot host the Manager virtual machine. You can convert some or all of these hosts to self-hosted engine nodes.

CHAPTER 8. REINSTALLING AN EXISTING HOST AS A SELF-HOSTED ENGINE NODE

You can convert an existing, standard host in a self-hosted engine environment to a self-hosted engine node capable of hosting the Manager virtual machine.

Procedure

1. Click **Compute** → **Hosts** and select the host.
2. Click **Management** → **Maintenance** and click **OK**.
3. Click **Installation** → **Reinstall**.
4. Click the **Hosted Engine** tab and select **DEPLOY** from the drop-down list.
5. Click **OK**.

The host is reinstalled with self-hosted engine configuration, and is flagged with a crown icon in the Administration Portal.

After reinstalling the hosts as self-hosted engine nodes, you can check the status of the new environment by running the following command on one of the nodes:

```
# hosted-engine --vm-status
```

If the new environment is running without issue, you can decommission the original Manager machine.