



# Red Hat Enterprise Virtualization 3.4

## Installation Guide

Installing Red Hat Enterprise Virtualization



# Red Hat Enterprise Virtualization 3.4 Installation Guide

---

Installing Red Hat Enterprise Virtualization

Andrew Burden  
aburden@redhat.com

Steve Gordon  
sgordon@redhat.com

Tim Hildred  
thildred@redhat.com

Cheryn Tan  
chetan@redhat.com

## Legal Notice

Copyright © 2015 Red Hat.

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

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, 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<sup>®</sup> is the registered trademark of Linus Torvalds in the United States and other countries.

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

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

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

Node.js<sup>®</sup> 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<sup>®</sup> 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 Enterprise Virtualization.

# Table of Contents

<b>PART I. INTRODUCTION</b>	<b>7</b>
<b>CHAPTER 1. INTRODUCTION</b>	<b>8</b>
1.1. WORKFLOW PROGRESS - SYSTEM REQUIREMENTS	8
1.2. RED HAT ENTERPRISE VIRTUALIZATION MANAGER REQUIREMENTS	8
1.2.1. Red Hat Enterprise Virtualization Hardware Requirements Overview	8
1.2.2. Red Hat Enterprise Virtualization Manager Hardware Requirements	9
1.2.3. Operating System Requirements	9
1.2.4. Red Hat Enterprise Virtualization Manager Client Requirements	10
1.2.5. Red Hat Enterprise Virtualization Manager Software Channels	11
1.3. HYPERVISOR REQUIREMENTS	11
1.3.1. Virtualization Host Hardware Requirements Overview	11
1.3.2. Virtualization Host CPU Requirements	12
1.3.3. Virtualization Host RAM Requirements	13
1.3.4. Virtualization Host Storage Requirements	13
1.3.5. Virtualization Host PCI Device Requirements	15
1.4. USER AUTHENTICATION	15
1.4.1. About Directory Services	15
1.4.2. Directory Services Support in Red Hat Enterprise Virtualization	15
1.5. FIREWALLS	17
1.5.1. Red Hat Enterprise Virtualization Manager Firewall Requirements	17
1.5.2. Virtualization Host Firewall Requirements	19
1.5.3. Directory Server Firewall Requirements	22
1.5.4. Database Server Firewall Requirements	22
1.6. SYSTEM ACCOUNTS	23
1.6.1. Red Hat Enterprise Virtualization Manager User Accounts	23
1.6.2. Red Hat Enterprise Virtualization Manager Groups	23
1.6.3. Virtualization Host User Accounts	23
1.6.4. Virtualization Host Groups	24
<b>PART II. INSTALLING RED HAT ENTERPRISE VIRTUALIZATION</b>	<b>25</b>
<b>CHAPTER 2. INSTALLING RED HAT ENTERPRISE VIRTUALIZATION</b>	<b>26</b>
2.1. WORKFLOW PROGRESS - INSTALLING RED HAT ENTERPRISE VIRTUALIZATION MANAGER	26
2.2. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION MANAGER	26
2.3. SUBSCRIBING TO THE REQUIRED CHANNELS	29
2.3.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager	29
2.3.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels Using RHN Classic	29
2.4. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION MANAGER	31
2.4.1. Configuring an Offline Repository for Red Hat Enterprise Virtualization Manager Installation	31
2.4.2. Installing the Red Hat Enterprise Virtualization Manager Packages	34
2.4.3. Configuring the Red Hat Enterprise Virtualization Manager	35
2.4.4. Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager	38
2.4.5. Configuring the Manager to Use a Manually Configured Local or Remote PostgreSQL Database	39
2.4.6. Connecting to the Administration Portal	40
2.4.7. Removing the Red Hat Enterprise Virtualization Manager	41
2.5. SPICE CLIENT	42
2.5.1. SPICE Features	42
<b>CHAPTER 3. THE SELF-HOSTED ENGINE</b>	<b>44</b>
3.1. ABOUT THE SELF-HOSTED ENGINE	44
3.2. LIMITATIONS OF THE SELF-HOSTED ENGINE	44

3.3. INSTALLING THE SELF-HOSTED ENGINE	44
3.4. CONFIGURING THE SELF-HOSTED ENGINE	45
3.5. INSTALLING ADDITIONAL HOSTS TO A SELF-HOSTED ENVIRONMENT	51
3.6. MAINTAINING THE SELF-HOSTED ENGINE	52
3.7. UPGRADING THE SELF-HOSTED ENGINE	53
3.8. UPGRADING ADDITIONAL HOSTS IN A SELF-HOSTED ENVIRONMENT	55
3.9. BACKING UP AND RESTORING A SELF-HOSTED ENVIRONMENT	56
3.9.1. Backing up the Self-Hosted Engine Manager Virtual Machine	56
3.9.2. Creating a New Self-Hosted Engine Environment to be Used as the Restored Environment	58
3.9.3. Restoring the Self-Hosted Engine Manager	63
3.9.4. Removing Non-Operational Hosts from a Restored Self-Hosted Engine Environment	69
3.9.5. Installing Additional Hosts to a Restored Self-Hosted Engine Environment	70
3.10. MIGRATING TO A SELF-HOSTED ENVIRONMENT	73
<b>CHAPTER 4. HISTORY AND REPORTS</b>	<b>81</b>
4.1. WORKFLOW PROGRESS - DATA COLLECTION SETUP AND REPORTS INSTALLATION	81
4.2. DATA COLLECTION SETUP AND REPORTS INSTALLATION OVERVIEW	81
4.3. INSTALLING AND CONFIGURING THE HISTORY DATABASE AND RED HAT ENTERPRISE VIRTUALIZATION MANAGER REPORTS	82
<b>CHAPTER 5. UPDATING THE RED HAT ENTERPRISE VIRTUALIZATION ENVIRONMENT</b>	<b>85</b>
5.1. UPGRADES BETWEEN MINOR RELEASES	85
5.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates	85
5.1.2. Updating the Red Hat Enterprise Virtualization Manager	86
5.1.3. Updating Red Hat Enterprise Virtualization Hypervisors	87
5.1.4. Updating Red Hat Enterprise Linux Virtualization Hosts	89
5.1.5. Updating the Red Hat Enterprise Virtualization Guest Tools	89
5.2. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION 3.4	90
5.2.1. Red Hat Enterprise Virtualization Manager 3.4 Upgrade Overview	90
5.2.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.4	90
5.2.3. Red Hat Enterprise Virtualization 3.4 Upgrade Considerations	91
5.2.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.4	92
5.3. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION 3.3	93
5.3.1. Red Hat Enterprise Virtualization Manager 3.3 Upgrade Overview	93
5.3.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3	94
5.3.3. Red Hat Enterprise Virtualization 3.3 Upgrade Considerations	95
5.3.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.3	96
5.4. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER 3.2	98
5.4.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2	98
5.5. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER 3.1	100
5.5.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.1	100
5.6. POST-UPGRADE TASKS	103
5.6.1. Changing the Cluster Compatibility Version	103
5.6.2. Changing the Data Center Compatibility Version	104
<b>PART III. INSTALLING HOSTS</b>	<b>106</b>
<b>CHAPTER 6. INTRODUCTION TO HOSTS</b>	<b>107</b>
6.1. WORKFLOW PROGRESS - INSTALLING VIRTUALIZATION HOSTS	107
6.2. INTRODUCTION TO VIRTUALIZATION HOSTS	107
<b>CHAPTER 7. RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR HOSTS</b>	<b>109</b>
7.1. RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR INSTALLATION OVERVIEW	109
7.2. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR DISK IMAGE	109

---

7.3. PREPARING INSTALLATION MEDIA	110
7.3.1. Preparing a USB Storage Device	110
7.3.2. Preparing USB Installation Media Using livecd-iso-to-disk	110
7.3.3. Preparing USB Installation Media Using dd	112
7.3.4. Preparing USB Installation Media Using dd on Linux Systems	113
7.3.5. Preparing USB Installation Media Using dd on Windows Systems	113
7.3.6. Preparing Optical Hypervisor Installation Media	115
7.4. INSTALLATION	116
7.4.1. Booting the Hypervisor from USB Installation Media	116
7.4.2. Booting the Hypervisor from Optical Installation Media	117
7.4.3. Starting the Installation Program	117
7.4.4. Hypervisor Menu Actions	120
7.4.5. Installing the Hypervisor	120
7.5. CONFIGURATION	123
7.5.1. Logging Into the Hypervisor	123
7.5.2. The Status Screen	124
7.5.3. The Network Screen	124
7.5.3.1. The Network Screen	124
7.5.3.2. Configuring the Host Name	124
7.5.3.3. Configuring Domain Name Servers	125
7.5.3.4. Configuring Network Time Protocol Servers	125
7.5.3.5. Configuring Network Interfaces	126
7.5.4. The Security Screen	127
7.5.5. The Keyboard Screen	127
7.5.6. The SNMP Screen	128
7.5.7. The CIM Screen	128
7.5.8. The Logging Screen	129
7.5.9. The Kdump Screen	129
7.5.10. The Remote Storage Screen	130
7.5.11. The Diagnostics Screen	131
7.5.12. The Performance Screen	131
7.5.13. The RHEV-M Screen	132
7.5.14. The Plugins Screen	133
7.5.15. The RHN Registration Screen	133
7.6. ADDING HYPERVISORS TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER	135
7.6.1. Using the Hypervisor	135
7.6.2. Approving a Hypervisor	135
7.7. MODIFYING THE RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR ISO	136
7.7.1. Introduction to Modifying the Red Hat Enterprise Virtualization Hypervisor ISO	136
7.7.2. Installing the edit-node Tool	136
7.7.3. Syntax of the edit-node Tool	137
7.7.4. Adding and Updating Packages	139
7.7.4.1. Creating a Local Repository	139
7.7.4.2. Example: Adding Packages to the Red Hat Enterprise Virtualization Hypervisor ISO File	140
7.7.4.3. Example: Updating Packages in the Red Hat Enterprise Virtualization Hypervisor ISO File	141
7.7.5. Modifying the Default ID of Users and Groups	141
7.7.5.1. Example: Modifying the Default ID of a User	141
7.7.5.2. Example: Modifying the Default ID of a Group	142
<b>CHAPTER 8. RED HAT ENTERPRISE LINUX HOSTS</b> .....	<b>143</b>
8.1. RED HAT ENTERPRISE LINUX HOSTS	143
8.2. HOST COMPATIBILITY MATRIX	143

---

<b>PART IV. BASIC SETUP</b>	<b>144</b>
<b>CHAPTER 9. CONFIGURING HOSTS</b>	<b>145</b>
9.1. INSTALLING RED HAT ENTERPRISE LINUX	145
9.2. SUBSCRIBING TO REQUIRED CHANNELS USING SUBSCRIPTION MANAGER	147
9.3. SUBSCRIBING TO REQUIRED CHANNELS USING RHN CLASSIC	148
9.4. CONFIGURING VIRTUALIZATION HOST FIREWALL	149
9.5. CONFIGURING VIRTUALIZATION HOST SUDO	151
9.6. CONFIGURING VIRTUALIZATION HOST SSH	151
9.7. ADDING A RED HAT ENTERPRISE LINUX HOST	152
9.8. EXPLANATION OF SETTINGS AND CONTROLS IN THE NEW HOST AND EDIT HOST WINDOWS	153
9.8.1. Host General Settings Explained	153
9.8.2. Host Power Management Settings Explained	155
9.8.3. SPM Priority Settings Explained	157
9.8.4. Host Console Settings Explained	157
<b>CHAPTER 10. CONFIGURING DATA CENTERS</b>	<b>159</b>
10.1. WORKFLOW PROGRESS - PLANNING YOUR DATA CENTER	159
10.2. PLANNING YOUR DATA CENTER	159
10.3. DATA CENTERS IN RED HAT ENTERPRISE VIRTUALIZATION	160
10.4. CREATING A NEW DATA CENTER	160
10.5. CHANGING THE DATA CENTER COMPATIBILITY VERSION	161
<b>CHAPTER 11. CONFIGURING CLUSTERS</b>	<b>163</b>
11.1. CLUSTERS IN RED HAT ENTERPRISE VIRTUALIZATION	163
11.2. CREATING A NEW CLUSTER	163
11.3. CHANGING THE CLUSTER COMPATIBILITY VERSION	164
<b>CHAPTER 12. CONFIGURING NETWORKING</b>	<b>166</b>
12.1. WORKFLOW PROGRESS - NETWORK SETUP	166
12.2. NETWORKING IN RED HAT ENTERPRISE VIRTUALIZATION	166
12.3. CREATING LOGICAL NETWORKS	168
12.3.1. Creating a New Logical Network in a Data Center or Cluster	168
12.4. EDITING LOGICAL NETWORKS	169
12.4.1. Editing Host Network Interfaces and Assigning Logical Networks to Hosts	169
12.4.2. Logical Network General Settings Explained	171
12.4.3. Editing a Logical Network	172
12.4.4. Explanation of Settings in the Manage Networks Window	173
12.4.5. Adding Multiple VLANs to a Single Network Interface Using Logical Networks	173
12.4.6. Multiple Gateways	175
12.4.7. Using the Networks Tab	175
12.5. EXTERNAL PROVIDER NETWORKS	176
12.5.1. Importing Networks From External Providers	176
12.5.2. Limitations to Using External Provider Networks	178
12.5.3. Configuring Subnets on External Provider Logical Networks	179
12.5.3.1. Configuring Subnets on External Provider Logical Networks	179
12.5.3.2. Adding Subnets to External Provider Logical Networks	179
12.5.3.3. Removing Subnets from External Provider Logical Networks	180
12.6. BONDING	180
12.6.1. Bonding Logic in Red Hat Enterprise Virtualization	180
12.6.2. Bonding Modes	181
12.6.3. Creating a Bond Device Using the Administration Portal	182
12.6.4. Example Uses of Custom Bonding Options with Host Interfaces	183
12.7. REMOVING LOGICAL NETWORKS	184



---

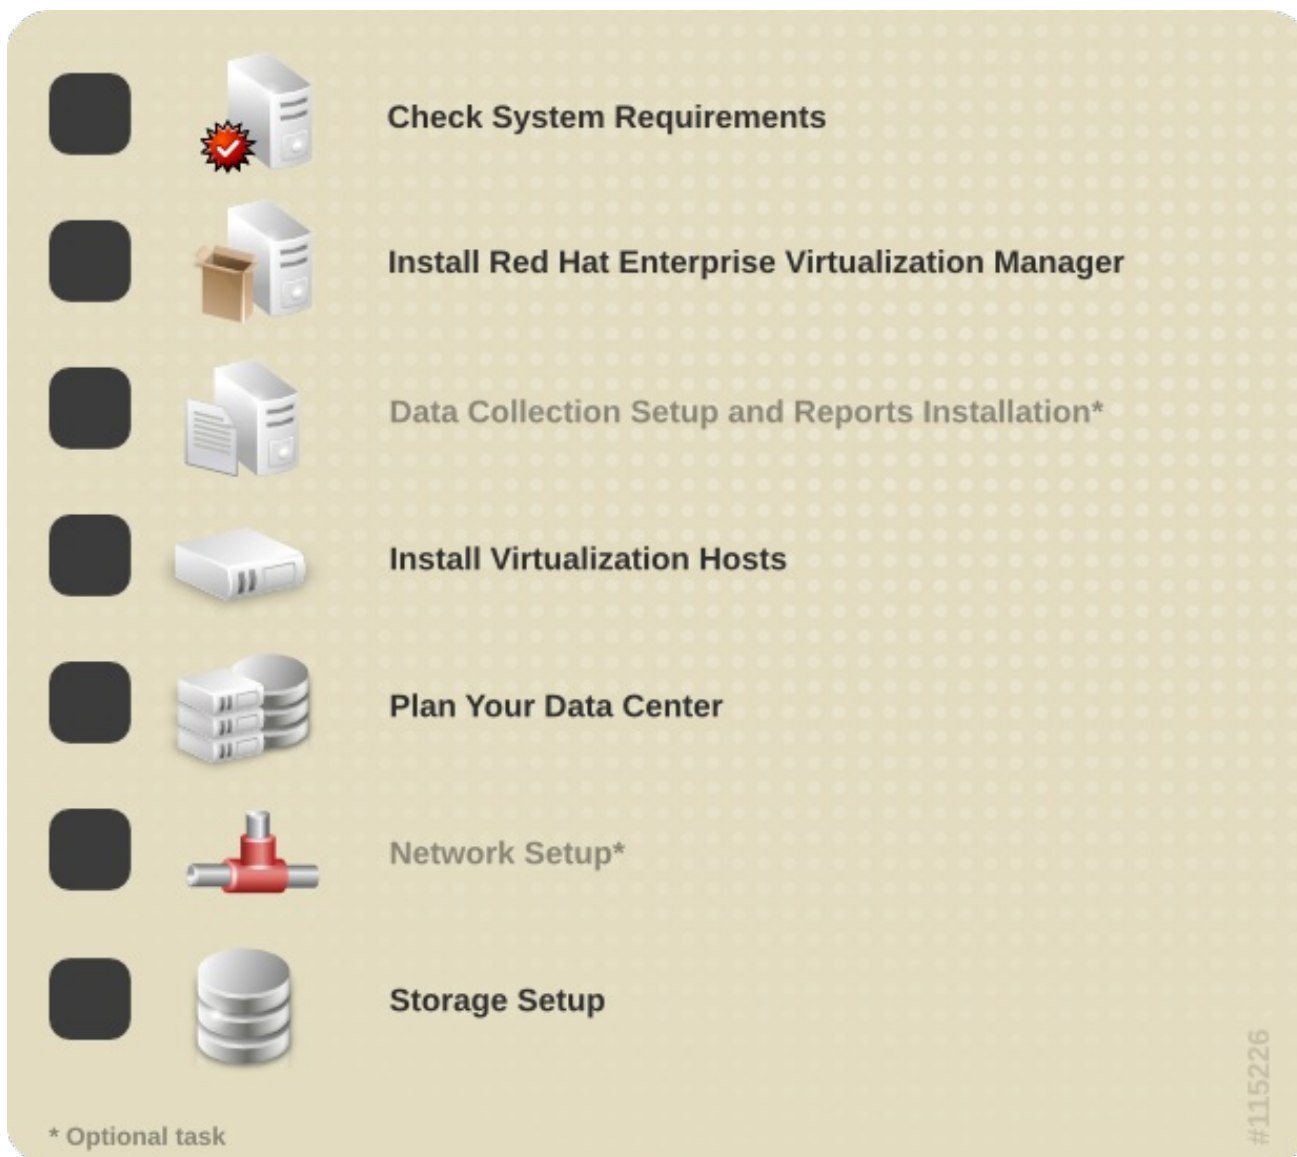
12.7.1. Removing a Logical Network	184
<b>CHAPTER 13. CONFIGURING STORAGE</b>	<b>186</b>
13.1. WORKFLOW PROGRESS - STORAGE SETUP	186
13.2. INTRODUCTION TO STORAGE IN RED HAT ENTERPRISE VIRTUALIZATION	186
13.3. PREPARING NFS STORAGE	188
13.4. ATTACHING NFS STORAGE	189
13.5. PREPARING PNFS STORAGE	191
13.6. ATTACHING PNFS STORAGE	191
13.7. ADDING ISCSI STORAGE	193
13.8. ADDING FCP STORAGE	195
13.9. PREPARING LOCAL STORAGE	197
13.10. ADDING LOCAL STORAGE	198
13.11. POSIX COMPLIANT FILE SYSTEM STORAGE IN RED HAT ENTERPRISE VIRTUALIZATION	199
13.12. ATTACHING POSIX COMPLIANT FILE SYSTEM STORAGE	199
13.13. ENABLING GLUSTER PROCESSES ON RED HAT STORAGE NODES	201
13.14. POPULATING THE ISO STORAGE DOMAIN	202
13.15. VIRTIO AND GUEST TOOL IMAGE FILES	202
13.16. UPLOADING THE VIRTIO AND GUEST TOOL IMAGE FILES TO AN ISO STORAGE DOMAIN	203
<b>CHAPTER 14. CONFIGURING LOGS</b>	<b>204</b>
14.1. RED HAT ENTERPRISE VIRTUALIZATION MANAGER INSTALLATION LOG FILES	204
14.2. RED HAT ENTERPRISE VIRTUALIZATION MANAGER LOG FILES	204
14.3. RED HAT ENTERPRISE VIRTUALIZATION HOST LOG FILES	205
14.4. SETTING UP A VIRTUALIZATION HOST LOGGING SERVER	205
14.5. THE LOGGING SCREEN	206
<b>PART V. ADVANCED SETUP</b>	<b>208</b>
<b>CHAPTER 15. PROXIES</b>	<b>209</b>
15.1. SPICE PROXY	209
15.1.1. SPICE Proxy Overview	209
15.1.2. SPICE Proxy Machine Setup	209
15.1.3. Turning on SPICE Proxy	210
15.1.4. Turning Off a SPICE Proxy	210
15.2. SQUID PROXY	211
15.2.1. Installing and Configuring a Squid Proxy	211
<b>APPENDIX A. REVISION HISTORY</b>	<b>216</b>



# PART I. INTRODUCTION

## CHAPTER 1. INTRODUCTION

### 1.1. WORKFLOW PROGRESS - SYSTEM REQUIREMENTS



[Report a bug](#)

## 1.2. RED HAT ENTERPRISE VIRTUALIZATION MANAGER REQUIREMENTS

### 1.2.1. Red Hat Enterprise Virtualization Hardware Requirements Overview

This section outlines the minimum hardware required to install, configure, and operate a Red Hat Enterprise Virtualization environment. To setup a Red Hat Enterprise Virtualization environment it is necessary to have, at least:

- one machine to act as the management server,
- one or more machines to act as virtualization hosts - at least two are required to support migration and power management,

- one or more machines to use as clients for accessing the Administration Portal.
- storage infrastructure provided by NFS, POSIX, iSCSI, SAN, or local storage.

The hardware required for each of these systems is further outlined in the following sections. The Red Hat Enterprise Virtualization environment also requires storage infrastructure that is accessible to the virtualization hosts. Storage infrastructure must be accessible using NFS, iSCSI, FC, or locally attached to virtualization hosts. The use of other POSIX compliant filesystems is also supported.

[Report a bug](#)

## 1.2.2. Red Hat Enterprise Virtualization Manager Hardware Requirements

The minimum and recommended hardware requirements outlined here are based on a typical small to medium sized installation. The exact requirements vary between deployments based on sizing and load. Please use these recommendations as a guide only.

### Minimum

- A dual core CPU.
- 4 GB of available system RAM that is not being consumed by existing processes.
- 25 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

### Recommended

- A quad core CPU or multiple dual core CPUs.
- 16 GB of system RAM.
- 50 GB of locally accessible, writeable, disk space.
- 1 Network Interface Card (NIC) with bandwidth of at least 1 Gbps.

The Red Hat Enterprise Virtualization Manager runs on Red Hat Enterprise Linux. To confirm whether or not specific hardware items are certified for use with Red Hat Enterprise Linux refer to <http://www.redhat.com/rhel/compatibility/hardware/>.

[Report a bug](#)

## 1.2.3. Operating System Requirements

The Red Hat Enterprise Virtualization Manager must run on Red Hat Enterprise Linux Server 6.5 or 6.6. You must install the operating system before installing the Red Hat Enterprise Virtualization Manager.

Moreover, the Red Hat Enterprise Virtualization Manager must be installed on a base installation of Red Hat Enterprise Linux. Do not install any additional packages after the base installation because they may cause dependency issues when attempting to install the packages required by the Manager.

**IMPORTANT**

See the [Red Hat Enterprise Linux 6 Security Guide](#) or the [Red Hat Enterprise Linux 7 Security Guide](#) for security hardening information for your Red Hat Enterprise Linux Servers.

[Report a bug](#)

### 1.2.4. Red Hat Enterprise Virtualization Manager Client Requirements

Use a client with a supported web browser to access the Administration Portal, and the User Portal. The portals support the following clients and browsers:

- Mozilla Firefox 17, and later, on Red Hat Enterprise Linux is required to access both portals.
- Internet Explorer 8, and later, on Microsoft Windows is required to access the User Portal. Use the desktop version, not the touchscreen version of Internet Explorer 10.
- Internet Explorer 9, and later, on Microsoft Windows is required to access the Administration Portal. Use the desktop version, not the touchscreen version of Internet Explorer 10.

Install a supported SPICE client to access virtual machine consoles. Supported SPICE clients are available on the following operating systems:

- Red Hat Enterprise Linux 5.8+ (i386, AMD64 and Intel 64)
- Red Hat Enterprise Linux 6.2+ (i386, AMD64 and Intel 64)
- Red Hat Enterprise Linux 6.5+ (i386, AMD64 and Intel 64)
- Windows XP
- Windows XP Embedded (XPe)
- Windows 7 (x86, AMD64 and Intel 64)
- Windows 8 (x86, AMD64 and Intel 64)
- Windows Embedded Standard 7
- Windows 2008/R2 (x86, AMD64 and Intel 64)
- Windows Embedded Standard 2009
- Red Hat Enterprise Virtualization Certified Linux-based thin clients

**NOTE**

Check the Red Hat Enterprise Virtualization Manager Release Notes to see which SPICE features your client supports.


When you access the portal(s) using Mozilla Firefox the SPICE client is provided by the `spice-xpi` package, which you must manually install using **yum**.

When you access the portal(s) using Internet Explorer the SPICE ActiveX control automatically downloads and installs.

[Report a bug](#)

## 1.2.5. Red Hat Enterprise Virtualization Manager Software Channels

To install Red Hat Enterprise Virtualization Manager, the system must be subscribed to the Red Hat Channels that deliver Red Hat Enterprise Virtualization and the channels that deliver Red Hat Enterprise Linux. These channels provide installation packages as well as updates.



### NOTE

See the *Red Hat Enterprise Virtualization Manager Release Notes* for specific channel names current to your system.

You must ensure that you have entitlements to the required channels listed here before proceeding with installation.

### Certificate-based Red Hat Network

- The **Red Hat Enterprise Linux Server** entitlement, provides Red Hat Enterprise Linux.
- The **Red Hat Enterprise Virtualization** entitlement, provides Red Hat Enterprise Virtualization Manager.
- The **Red Hat JBoss Enterprise Application Platform** entitlement, provides the supported release of the application platform on which the Manager runs.

### Red Hat Network Classic

- The **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86\_64)** channel, also referred to as **rhel-x86\_64-server-6**, provides Red Hat Enterprise Linux 6 Server. The Channel Entitlement name for this channel is **Red Hat Enterprise Linux Server (v. 6)**.
- The **RHEL Server Supplementary (v. 6 64-bit x86\_64)** channel, also referred to as **rhel-x86\_64-server-supplementary-6**, provides the virtio-win package. The virtio-win package provides the Windows VirtIO drivers for use in virtual machines. The Channel Entitlement Name for the supplementary channel is **Red Hat Enterprise Linux Server Supplementary (v. 6)**.
- The **Red Hat Enterprise Virtualization Manager (v3.4 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.4**, provides Red Hat Enterprise Virtualization Manager. The Channel Entitlement Name for this channel is **Red Hat Enterprise Virtualization Manager (v3)**.
- The **Red Hat JBoss EAP (v 6) for 6Server x86\_64** channel, also referred to as **jbappplatform-6-x86\_64-server-6-rpm**, provides the supported release of the application platform on which the Manager runs. The Channel Entitlement Name for this channel is **Red Hat JBoss Enterprise Application Platform (v 4, zip format)**.

[Report a bug](#)

## 1.3. HYPERVISOR REQUIREMENTS

### 1.3.1. Virtualization Host Hardware Requirements Overview

Red Hat Enterprise Virtualization Hypervisors and Red Hat Enterprise Linux Hosts have a number of hardware requirements and supported limits.

[Report a bug](#)

### 1.3.2. Virtualization Host CPU Requirements

Red Hat Enterprise Virtualization supports the use of these CPU models in virtualization hosts:

- AMD Opteron G1
- AMD Opteron G2
- AMD Opteron G3
- AMD Opteron G4
- AMD Opteron G5
- Intel Conroe
- Intel Penryn
- Intel Nehalem
- Intel Westmere
- Intel Sandybridge
- Intel Haswell

All CPUs must have support for the Intel® 64 or AMD64 CPU extensions, and the AMD-V™ or Intel VT® hardware virtualization extensions enabled. Support for the **No eXecute** flag (NX) is also required. To check that your processor supports the required flags, and that they are enabled:

1. At the Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor boot screen, press any key and select the **Boot** or **Boot with serial console** entry from the list.
2. Press **Tab** to edit the kernel parameters for the selected option.
3. Ensure there is a **Space** after the last kernel parameter listed, and append the **rescue** parameter.
4. Press **Enter** to boot into rescue mode.
5. At the prompt which appears, determine that your processor has the required extensions and that they are enabled by running this command:

```
# grep -E 'svm|vmx' /proc/cpuinfo | grep nx
```

If any output is shown, then the processor is hardware virtualization capable. If no output is shown, then it is still possible that your processor supports hardware virtualization. In some circumstances manufacturers disable the virtualization extensions in the BIOS. If you believe this to be the case, consult the system's BIOS and the motherboard manual provided by the manufacturer.

6. As an additional check, verify that the **kvm** modules are loaded in the kernel:



```
# lsmod | grep kvm
```

If the output includes **kvm\_intel** or **kvm\_amd** then the **kvm** hardware virtualization modules are loaded and your system meets requirements.

[Report a bug](#)

### 1.3.3. Virtualization Host RAM Requirements

It is recommended that virtualization hosts have at least 2 GB of RAM. The amount of RAM required varies depending on:

- guest operating system requirements,
- guest application requirements, and
- memory activity and usage of guests.

The fact that KVM is able to over-commit physical RAM for virtualized guests must also be taken into account. This allows for provisioning of guests with RAM requirements greater than what is physically present, on the basis that the guests are not all concurrently at peak load. KVM does this by only allocating RAM for guests as required and shifting underutilized guests into swap.

A maximum of 2 TB of RAM per virtualization host is currently supported.

[Report a bug](#)

### 1.3.4. Virtualization Host Storage Requirements

Virtualization hosts require local storage to store configuration, logs, kernel dumps, and for use as swap space. The minimum storage requirements of the Red Hat Enterprise Virtualization Hypervisor are documented in this section. The storage requirements for Red Hat Enterprise Linux hosts vary based on the amount of disk space used by their existing configuration but are expected to be greater than those of the Red Hat Enterprise Virtualization Hypervisor.

It is recommended that each virtualization host has at least 2 GB of internal storage. The minimum supported internal storage for each Hypervisor is the total of that required to provision the following partitions:

- The root partitions require at least 512 MB of storage.
- The configuration partition requires at least 8 MB of storage.
- The recommended minimum size of the logging partition is 2048 MB.
- The data partition requires at least 256 MB of storage. Use of a smaller data partition may prevent future upgrades of the Hypervisor from the Red Hat Enterprise Virtualization Manager. By default all disk space remaining after allocation of swap space will be allocated to the data partition.
- The swap partition requires at least 8 MB of storage. The recommended size of the swap partition varies depending on both the system the Hypervisor is being installed upon and the anticipated level of overcommit for the environment. Overcommit allows the Red Hat Enterprise Virtualization environment to present more RAM to guests than is actually physically present. The default overcommit ratio is **0.5**.

The recommended size of the swap partition can be determined by:

- Multiplying the amount of system RAM by the expected overcommit ratio, and adding
- 2 GB of swap space for systems with 4 GB of RAM or less, or
- 4 GB of swap space for systems with between 4 GB and 16 GB of RAM, or
- 8 GB of swap space for systems with between 16 GB and 64 GB of RAM, or
- 16 GB of swap space for systems with between 64 GB and 256 GB of RAM.

### Example 1.1. Calculating Swap Partition Size

For a system with 8 GB of RAM this means the formula for determining the amount of swap space to allocate is:

$$(8 \text{ GB} \times 0.5) + 4 \text{ GB} = 8 \text{ GB}$$

### IMPORTANT

By default the Red Hat Enterprise Virtualization Hypervisor defines a swap partition sized using the recommended formula. An overcommit ratio of **0.5** is used for this calculation. For some systems the result of this calculation may be a swap partition that requires more free disk space than is available at installation. Where this is the case Hypervisor installation will fail.

If you encounter this issue, manually set the sizes for the Hypervisor disk partitions using the ***storage\_vol*** boot parameter.

### Example 1.2. Manually Setting Swap Partition Size

In this example the ***storage\_vol*** boot parameter is used to set a swap partition size of 4096 MB. Note that no sizes are specified for the other partitions, allowing the Hypervisor to use the default sizes.

```
storage_vol=:4096:::
```

### IMPORTANT

The Red Hat Enterprise Virtualization Hypervisor does not support installation on **fakeraid** devices. Where a **fakeraid** device is present it must be reconfigured such that it no longer runs in RAID mode.

1. Access the RAID controller's BIOS and remove all logical drives from it.
2. Change controller mode to be non-RAID. This may be referred to as compatibility or JBOD mode.

Access the manufacturer provided documentation for further information related to the specific device in use.

---

[Report a bug](#)

### 1.3.5. Virtualization Host PCI Device Requirements

Virtualization hosts must have at least one network interface with a minimum bandwidth of 1 Gbps. It is recommended that each virtualization host have two network interfaces with a minimum bandwidth of 1 Gbps to support network intensive activity, including virtual machine migration.

[Report a bug](#)

## 1.4. USER AUTHENTICATION

### 1.4.1. About Directory Services

The term directory service refers to the collection of software, hardware, and processes that store information about an enterprise, subscribers, or both, and make that information available to users. A directory service consists of at least one directory server and at least one directory client program. Client programs can access names, phone numbers, addresses, and other data stored in the directory service.

[Report a bug](#)

### 1.4.2. Directory Services Support in Red Hat Enterprise Virtualization

During installation Red Hat Enterprise Virtualization Manager creates its own internal administration user, **admin**. This account is intended for use when initially configuring the environment, and for troubleshooting. To add other users to Red Hat Enterprise Virtualization you must attach a directory server to the Manager using the Domain Management Tool, **engine-manage-domains**.

Once at least one directory server has been attached to the Manager, you can add users that exist in the directory server and assign roles to them using the Administration Portal. Users can be identified by their User Principal Name (UPN) of the form **user@domain**. Attachment of more than one directory server to the Manager is also supported.

The directory servers supported for use with Red Hat Enterprise Virtualization 3.4 are:

- Active Directory
- Identity Management (IdM)
- Red Hat Directory Server 9 (RHDS 9)
- OpenLDAP

You must ensure that the correct DNS records exist for your directory server. In particular you must ensure that the DNS records for the directory server include:

- A valid pointer record (PTR) for the directory server's reverse look-up address.
- A valid service record (SRV) for LDAP over TCP port **389**.
- A valid service record (SRV) for Kerberos over TCP port **88**.
- A valid service record (SRV) for Kerberos over UDP port **88**.

If these records do not exist in DNS then you cannot add the domain to the Red Hat Enterprise Virtualization Manager configuration using **engine-manage-domains**.

For more detailed information on installing and configuring a supported directory server, see the vendor's documentation:

- Active Directory - <http://technet.microsoft.com/en-us/windowsserver/dd448614>.
- Identity Management (IdM) - [http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Enterprise\\_Linux/6/html/Identity\\_Management\\_Guide/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Identity_Management_Guide/index.html)
- Red Hat Directory Server (RHDS) - [http://docs.redhat.com/docs/en-US/Red\\_Hat\\_Directory\\_Server/index.html](http://docs.redhat.com/docs/en-US/Red_Hat_Directory_Server/index.html)
- OpenLDAP - <http://www.openldap.org/doc/>



### IMPORTANT

A user must be created in the directory server specifically for use as the Red Hat Enterprise Virtualization administrative user. Do *not* use the administrative user for the directory server as the Red Hat Enterprise Virtualization administrative user.



### IMPORTANT

It is not possible to install Red Hat Enterprise Virtualization Manager (rhev) and IdM (ipa-server) on the same system. IdM is incompatible with the `mod_ssl` package, which is required by Red Hat Enterprise Virtualization Manager.



### IMPORTANT

If you are using Active Directory as your directory server, and you want to use **sysprep** in the creation of Templates and Virtual Machines, then the Red Hat Enterprise Virtualization administrative user must be delegated control over the Domain to:

- **Join a computer to the domain**
- **Modify the membership of a group**

For information on creation of user accounts in Active Directory, see <http://technet.microsoft.com/en-us/library/cc732336.aspx>.

For information on delegation of control in Active Directory, see <http://technet.microsoft.com/en-us/library/cc732524.aspx>.

**NOTE**

Red Hat Enterprise Virtualization Manager uses Kerberos to authenticate with directory servers. RHDS does not provide native support for Kerberos. If you are using RHDS as your directory server then you must ensure that the directory server is made a service within a valid Kerberos domain. To do this you must perform these steps while referring to the relevant directory server documentation:

- Configure the **memberOf** plug-in for RHDS to allow group membership. In particular ensure that the value of the **memberofgroupattr** attribute of the **memberOf** plug-in is set to **uniqueMember**. In **OpenLDAP**, the **memberOf** functionality is not called a "plugin". It is called an "overlay" and requires no configuration after installation.

Consult the Red Hat Directory Server 9.0 *Plug-in Guide* for more information on configuring the **memberOf** plug-in.

- Define the directory server as a service of the form **ldap/hostname@REALMNAME** in the Kerberos realm. Replace *hostname* with the fully qualified domain name associated with the directory server and **REALMNAME** with the fully qualified Kerberos realm name. The Kerberos realm name must be specified in capital letters.
- Generate a **keytab** file for the directory server in the Kerberos realm. The **keytab** file contains pairs of Kerberos principals and their associated encrypted keys. These keys allow the directory server to authenticate itself with the Kerberos realm.

Consult the documentation for your Kerberos principle for more information on generating a **keytab** file.

- Install the **keytab** file on the directory server. Then configure RHDS to recognize the **keytab** file and accept Kerberos authentication using GSSAPI.

Consult the Red Hat Directory Server 9.0 *Administration Guide* for more information on configuring RHDS to use an external **keytab** file.

- Test the configuration on the directory server by using the **kinit** command to authenticate as a user defined in the Kerberos realm. Once authenticated run the **ldapsearch** command against the directory server. Use the **-Y GSSAPI** parameters to ensure the use of Kerberos for authentication.

[Report a bug](#)

## 1.5. FIREWALLS

### 1.5.1. Red Hat Enterprise Virtualization Manager Firewall Requirements

The Red Hat Enterprise Virtualization Manager requires that a number of ports be opened to allow network traffic through the system's firewall. The **engine-setup** script is able to configure the firewall automatically, but this overwrites any pre-existing firewall configuration.

Where an existing firewall configuration exists, you must manually insert the firewall rules required by the Manager instead. The **engine-setup** command saves a list of the **iptables** rules required in the **/usr/share/ovirt-engine/conf/iptables.example** file.

The firewall configuration documented here assumes a default configuration. Where non-default HTTP and HTTPS ports are chosen during installation adjust the firewall rules to allow network traffic on the ports that were selected - not the default ports (**80** and **443**) listed here.

**Table 1.1. Red Hat Enterprise Virtualization Manager Firewall Requirements**

Port(s)	Protocol	Source	Destination	Purpose
-	ICMP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisor(s)</li> <li>Red Hat Enterprise Linux host(s)</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	When registering to the Red Hat Enterprise Virtualization Manager, virtualization hosts send an ICMP ping request to the Manager to confirm that it is online.
22	TCP	<ul style="list-style-type: none"> <li>System(s) used for maintenance of the Manager including backend configuration, and software upgrades.</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	SSH (optional)
80, 443	TCP	<ul style="list-style-type: none"> <li>Administration Portal clients</li> <li>User Portal clients</li> <li>Red Hat Enterprise Virtualization Hypervisor(s)</li> <li>Red Hat Enterprise Linux host(s)</li> <li>REST API clients</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	Provides HTTP and HTTPS access to the Manager.

**IMPORTANT**

In environments where the Red Hat Enterprise Virtualization Manager is also required to export NFS storage, such as an ISO Storage Domain, additional ports must be allowed through the firewall. Grant firewall exceptions for the ports applicable to the version of NFS in use:

**NFSv4**

- TCP port **2049** for NFS.

**NFSv3**

- TCP and UDP port **2049** for NFS.
- TCP and UDP port **111** (**rpcbind/sunrpc**).
- TCP and UDP port specified with **MOUNTD\_PORT="port"**
- TCP and UDP port specified with **STATD\_PORT="port"**
- TCP port specified with **LOCKD\_TCPPOINT="port"**
- UDP port specified with **LOCKD\_UDPOINT="port"**

The **MOUNTD\_PORT**, **STATD\_PORT**, **LOCKD\_TCPPOINT**, and **LOCKD\_UDPOINT** ports are configured in the **/etc/sysconfig/nfs** file.

[Report a bug](#)

## 1.5.2. Virtualization Host Firewall Requirements

Red Hat Enterprise Linux hosts and Red Hat Enterprise Virtualization Hypervisors require a number of ports to be opened to allow network traffic through the system's firewall. In the case of the Red Hat Enterprise Virtualization Hypervisor these firewall rules are configured automatically. For Red Hat Enterprise Linux hosts however it is necessary to manually configure the firewall.

**Table 1.2. Virtualization Host Firewall Requirements**

Port(s)	Protocol	Source	Destination	Purpose
22	TCP	<ul style="list-style-type: none"> <li>• Red Hat Enterprise Virtualization Manager</li> </ul>	<ul style="list-style-type: none"> <li>• Red Hat Enterprise Virtualization Hypervisors</li> <li>• Red Hat Enterprise Linux hosts</li> </ul>	Secure Shell (SSH) access.

Port(s)	Protocol	Source	Destination	Purpose
161	UDP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Linux Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	Simple network management protocol (SNMP).
5900 - 6923	TCP	<ul style="list-style-type: none"> <li>Administration Portal clients</li> <li>User Portal clients</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	Remote guest console access via VNC and SPICE. These ports must be open to facilitate client access to virtual machines.
5989	TCP, UDP	<ul style="list-style-type: none"> <li>Common Information Model Object Manager (CIMOM)</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	Used by Common Information Model Object Managers (CIMOM) to monitor virtual machines running on the virtualization host. To use a CIMOM to monitor the virtual machines in your virtualization environment then you must ensure that this port is open.



Port(s)	Protocol	Source	Destination	Purpose
16514	TCP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	Virtual machine migration using <b>libvirt</b> .
49152 - 49216	TCP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Linux Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Linux Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	Virtual machine migration and fencing using VDSM. These ports must be open facilitate both automated and manually initiated migration of virtual machines.
54321	TCP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Hypervisors</li> <li>Red Hat Enterprise Linux hosts</li> </ul>	VDSM communications with the Manager and other virtualization hosts.

### Example 1.3. Option Name: IPTablesConfig

Recommended (default) values: Automatically generated by vdsd bootstrap script

```
*filter
:INPUT ACCEPT [0:0]
:FORWARD ACCEPT [0:0]
:OUTPUT ACCEPT [0:0]
```

```

-A INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
-A INPUT -p icmp -j ACCEPT
-A INPUT -i lo -j ACCEPT
# vdsm
-A INPUT -p tcp --dport 54321 -j ACCEPT
# libvirt tls
-A INPUT -p tcp --dport 16514 -j ACCEPT
# SSH
-A INPUT -p tcp --dport 22 -j ACCEPT
# guest consoles
-A INPUT -p tcp -m multiport --dports 5900:6923 -j ACCEPT
# migration
-A INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
# snmp
-A INPUT -p udp --dport 161 -j ACCEPT
# Reject any other input traffic
-A INPUT -j REJECT --reject-with icmp-host-prohibited
-A FORWARD -m physdev ! --physdev-is-bridged -j REJECT --reject-with icmp-host-prohibited
COMMIT

```

[Report a bug](#)

### 1.5.3. Directory Server Firewall Requirements

Red Hat Enterprise Virtualization requires a directory server to support user authentication. A number of ports must be opened in the directory server's firewall to support GSS-API authentication as used by the Red Hat Enterprise Virtualization Manager.

**Table 1.3. Host Firewall Requirements**

Port(s)	Protocol	Source	Destination	Purpose
88, 464	TCP, UDP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	<ul style="list-style-type: none"> <li>Directory server</li> </ul>	Kerberos authentication.
389, 636	TCP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	<ul style="list-style-type: none"> <li>Directory server</li> </ul>	Lightweight Directory Access Protocol (LDAP) and LDAP over SSL.

[Report a bug](#)

### 1.5.4. Database Server Firewall Requirements

Red Hat Enterprise Virtualization supports the use of a remote database server. If you plan to use a remote database server with Red Hat Enterprise Virtualization then you must ensure that the remote database server allows connections from the Manager.

**Table 1.4. Host Firewall Requirements**

Port(s)	Protocol	Source	Destination	Purpose
5432	TCP, UDP	<ul style="list-style-type: none"> <li>Red Hat Enterprise Virtualization Manager</li> </ul>	<ul style="list-style-type: none"> <li>PostgreSQL database server</li> </ul>	Default port for PostgreSQL database connections.

If you plan to use a local database server on the Manager itself, which is the default option provided during installation, then no additional firewall rules are required.

[Report a bug](#)

## 1.6. SYSTEM ACCOUNTS

### 1.6.1. Red Hat Enterprise Virtualization Manager User Accounts

When the `rhev` package is installed, a number of user accounts are created to support Red Hat Enterprise Virtualization. The user accounts created as a result of `rhev` package installation are as follows. The default user identifier (UID) for each account is also provided:

- The **`vds`** user (UID **36**). Required for support tools that mount and access NFS storage domains.
- The **`ovirt`** user (UID **108**). Owner of the **`ovirt-engine`** Red Hat JBoss Enterprise Application Platform instance.

[Report a bug](#)

### 1.6.2. Red Hat Enterprise Virtualization Manager Groups

When the `rhev` package is installed, a number of user groups are created. The user groups created as a result of `rhev` package installation are as follows. The default group identifier (GID) for each group is also listed:

- The **`kvm`** group (GID **36**). Group members include:
  - The **`vds`** user.
- The **`ovirt`** group (GID **108**). Group members include:
  - The **`ovirt`** user.

[Report a bug](#)

### 1.6.3. Virtualization Host User Accounts

When the `vds` and `qemu-kvm-rhev` packages are installed, a number of user accounts are created. These are the user accounts that are created on the virtualization host as a result of `vds` and `qemu-kvm-rhev` package installation. The default user identifier (UID) for each entry is also listed:

- The **vds** user (UID **36**).
- The **qemu** user (UID **107**).
- The **sanlock** user (UID **179**).

In addition Red Hat Enterprise Virtualization Hypervisor hosts define a **admin** user (UID **500**). This **admin** user is not created on Red Hat Enterprise Linux virtualization hosts. The **admin** user is created with the required permissions to run commands as the **root** user using the **sudo** command. The **vds** user which is present on both types of virtualization hosts is also given access to the **sudo** command.



### IMPORTANT

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vds** user however is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system then a conflict will arise during installation of the `vds` and `qemu-kvm-rhev` packages.

[Report a bug](#)

## 1.6.4. Virtualization Host Groups

When the `vds` and `qemu-kvm-rhev` packages are installed, a number of user groups are created. These are the groups that are created on the virtualization host as a result of `vds` and `qemu-kvm-rhev` package installation. The default group identifier (GID) for each entry is also listed:

- The **kvm** group (GID **36**). Group members include:
  - The **qemu** user.
  - The **sanlock** user.
- The **qemu** group (GID **107**). Group members include:
  - The **vds** user.
  - The **sanlock** user.



### IMPORTANT

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vds** user however is fixed to a UID of **36** and the **kvm** group is fixed to a GID of **36**.

If UID **36** or GID **36** is already used by another account on the system then a conflict will arise during installation of the `vds` and `qemu-kvm-rhev` packages.

[Report a bug](#)

## PART II. INSTALLING RED HAT ENTERPRISE VIRTUALIZATION

## CHAPTER 2. INSTALLING RED HAT ENTERPRISE VIRTUALIZATION

### 2.1. WORKFLOW PROGRESS - INSTALLING RED HAT ENTERPRISE VIRTUALIZATION MANAGER




[Report a bug](#)

### 2.2. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION MANAGER

#### Overview

The Red Hat Enterprise Virtualization Manager can be installed under one of two arrangements - a standard setup in which the Manager is installed on an independent physical machine or virtual machine, or a self-hosted engine setup in which the Manager runs on a virtual machine that the Manager itself controls.



## IMPORTANT

While the prerequisites for and basic configuration of the Red Hat Enterprise Virtualization Manager itself are the same for both standard and self-hosted engine setups, the process for setting up a self-hosted engine is different from that of a standard setup.

### Prerequisites

Before installing the Red Hat Virtualization Manager, you must ensure that you meet all the prerequisites. To complete installation of the Red Hat Enterprise Virtualization Manager successfully, you must also be able to determine:

1. The ports to be used for HTTP and HTTPS communication. The default ports are **80** and **443** respectively.
2. The fully qualified domain name (FQDN) of the system on which the Manager is to be installed.
3. The password you use to secure the Red Hat Enterprise Virtualization administration account.
4. The location of the database server to be used. You can use the setup script to install and configure a local database server or use an existing remote database server. To use a remote database server you must know:
  - The host name of the system on which the remote database server exists.
  - The port on which the remote database server is listening.
  - That the **uuid-oss** extension had been loaded by the remote database server.

You must also know the user name and password of a user that is known to the remote database server. The user must have permission to create databases in PostgreSQL.

5. The organization name to use when creating the Manager's security certificates.
6. The storage type to be used for the initial data center attached to the Manager. The default is NFS.
7. The path to use for the ISO share, if the Manager is being configured to provide one. The display name, which will be used to label the domain in the Red Hat Enterprise Virtualization Manager also needs to be provided.
8. The firewall rules, if any, present on the system that need to be integrated with the rules required for the Manager to function.

### Configuration

Before installation is completed the values selected are displayed for confirmation. Once the values have been confirmed they are applied and the Red Hat Enterprise Virtualization Manager is ready for use.

#### Example 2.1. Completed Installation

```
--== CONFIGURATION PREVIEW ==--

Engine database name           : engine
Engine database secured connection : False
Engine database host           : localhost
```

```

Engine database user name      : engine
Engine database host name validation : False
Engine database port          : 5432
NFS setup                     : True
PKI organization              : Your Org
Application mode              : both
Firewall manager              : iptables
Update Firewall               : True
Configure WebSocket Proxy     : True
Host FQDN                     : Your Manager's FQDN
NFS export ACL                : 0.0.0.0/0.0.0.0(rw)
NFS mount point               : /var/lib/exports/iso
Datacenter storage type       : nfs
Configure local Engine database : True
Set application as default page : True
Configure Apache SSL          : True

```

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

## NOTE

Automated installations are created by providing **engine-setup** with an answer file. An answer file contains answers to the questions asked by the setup command.

- To create an answer file, use the ***--generate-answer*** parameter to specify a path and file name with which to create the answer file. When this option is specified, the **engine-setup** command records your answers to the questions in the setup process to the answer file.

```
# engine-setup --generate-answer=[ANSWER_FILE]
```

- To use an answer file for a new installation, use the ***--config-append*** parameter to specify the path and file name of the answer file to be used. The **engine-setup** command will use the answers stored in the file to complete the installation.

```
# engine-setup --config-append=[ANSWER_FILE]
```

Run **engine-setup --help** for a full list of parameters.

## NOTE

Offline installation requires the creation of a software repository local to your Red Hat Enterprise Virtualization environment. This software repository must contain all of the packages required to install Red Hat Enterprise Virtualization Manager, Red Hat Enterprise Linux virtualization hosts, and Red Hat Enterprise Linux virtual machines. To create such a repository, see the *Red Hat Enterprise Virtualization Manager Offline Installation* technical brief, available at <https://access.redhat.com/articles/216983>.

[Report a bug](#)



## 2.3. SUBSCRIBING TO THE REQUIRED CHANNELS

### 2.3.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager

#### Summary

Before you can install the Red Hat Enterprise Virtualization Manager, you must register the system on which the Red Hat Enterprise Virtualization Manager will be installed with the Red Hat Network and subscribe to the required channels.

#### Procedure 2.1. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager

1. **Register the System with Subscription Manager**

Run the following command and enter your Red Hat Network user name and password to register the system with the Red Hat Network:

```
# subscription-manager register
```

2. **Identify Available Entitlement Pools**

Run the following commands to find entitlement pools containing the channels required to install the Red Hat Enterprise Virtualization Manager:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Linux Server"
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Virtualization"
```

3. **Attach Entitlement Pools to the System**

Use the pool identifiers located in the previous step to attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the system. Run the following command to attach the entitlements:

```
# subscription-manager attach --pool=[POOLID]
```

4. **Enable the Required Channels**

Run the following commands to enable the channels required to install Red Hat Enterprise Virtualization:


```
# yum-config-manager --enable rhel-6-server-rpms
# yum-config-manager --enable rhel-6-server-supplementary-rpms
# yum-config-manager --enable rhel-6-server-rhev-3.4-rpms
# yum-config-manager --enable jboss-eap-6-for-rhel-6-server-rpms
```

#### Result

You have registered the system with Red Hat Network and subscribed to the channels required to install the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

### 2.3.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels Using RHN Classic

**NOTE**

See the *Red Hat Enterprise Virtualization Manager Release Notes* for channel names specific to your system.

**Summary**

To install Red Hat Enterprise Virtualization Manager you must first register the target system to Red Hat Network and subscribe to the required software channels.

**Procedure 2.2. Subscribing to the Red Hat Enterprise Virtualization Manager Channels using RHN Classic**

1. Run the **rhn\_register** command to register the system with Red Hat Network. To complete registration successfully you must supply your Red Hat Network user name and password. Follow the on-screen prompts to complete registration of the system.

```
# rhn_register
```

2. **Subscribe to Required Channels**

You must subscribe the system to the required channels using either the web interface to Red Hat Network or the command line **rhn-channel** command.

- o **Using the rhn-channel Command**

Run the **rhn-channel** command to subscribe the system to each of the required channels. Run the following commands:

```
# rhn-channel --add --channel=rhel-x86_64-server-6
# rhn-channel --add --channel=rhel-x86_64-server-supplementary-6
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.4
# rhn-channel --add --channel=jbappplatform-6-x86_64-server-6-rpm
```

**IMPORTANT**

If you are not the administrator for the machine as defined in Red Hat Network, or the machine is not registered to Red Hat Network, then use of the **rhn-channel** command results in an error:

```
Error communicating with server. The message was:
Error Class Code: 37
Error Class Info: You are not allowed to perform administrative tasks on
this system.
Explanation:
  An error has occurred while processing your request. If this problem
  persists please enter a bug report at bugzilla.redhat.com.
  If you choose to submit the bug report, please be sure to include
  details of what you were trying to do when this error occurred and
  details on how to reproduce this problem.
```

If you encounter this error when using **rhn-channel**, you must use the web user interface to add the channel.

- o **Using the Web Interface to Red Hat Network**

To add a channel subscription to a system from the web interface:

- a. Log on to Red Hat Network (<http://rhn.redhat.com>).
- b. Move the mouse cursor over the **Subscriptions** link at the top of the screen, and then click the **Registered Systems** link in the menu that appears.
- c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.
- d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.
- e. Select the channels to be added from the list presented on the screen. Red Hat Enterprise Virtualization Manager requires:
  - The **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86\_64)** channel. This channel is located under the **Release Channels for Red Hat Enterprise Linux 6 for x86\_64** expandable menu.
  - The **RHEL Server Supplementary (v. 6 64-bit x86\_64)** channel. This channel is located under the **Release Channels for Red Hat Enterprise Linux 6 for x86\_64** expandable menu.
  - The **Red Hat Enterprise Virtualization Manager (v.3.4 x86\_64)** channel. This channel is located under the **Additional Services Channels for Red Hat Enterprise Linux 6 for x86\_64** expandable menu.
  - The **Red Hat JBoss EAP (v 6) for 6Server x86\_64** channel. This channel is located under the **Additional Services Channels for Red Hat Enterprise Linux 6 for x86\_64** expandable menu.
- f. Click the **Change Subscription** button to finalize the change.

## Result

The system is now registered with Red Hat Network and subscribed to the channels required for Red Hat Enterprise Virtualization Manager installation.

[Report a bug](#)

## 2.4. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION MANAGER

### 2.4.1. Configuring an Offline Repository for Red Hat Enterprise Virtualization Manager Installation

This task documents the creation of an offline repository containing all packages needed to install a Red Hat Enterprise Virtualization environment. Follow these steps to create a repository you can use to install Red Hat Enterprise Virtualization components on systems without a direct connection to Red Hat Network.

#### 1. Install Red Hat Enterprise Linux

Install Red Hat Enterprise Linux 6 Server on a system that has access to Red Hat Network. This system downloads all required packages, and distribute them to your offline system(s).



## IMPORTANT

Ensure that the system used has a large amount of free disk space available. This procedure downloads a large number of packages, and requires up to 1.5 GB of free disk space.

### 2. Register Red Hat Enterprise Linux

Register the system with Red Hat Network (RHN) using either Subscription Manager or RHN Classic.

- **Subscription Manager**

Use the **subscription-manager** command as **root** with the **register** parameter.

```
# subscription-manager register
```

- **RHN Classic**

Use the **rhn\_register** command as **root**.

```
# rhn_register
```

### 3. Add required channel subscriptions

Subscribe the system for all channels listed in the *Red Hat Enterprise Virtualization - Installation Guide*.

- **Subscription Manager**

- **RHN Classic**

### 4. Configure File Transfer Protocol (FTP) access

Servers that are not connected to the Internet can access the software repository using File Transfer Protocol (FTP). To create the FTP repository you must install and configure vsftpd, while logged in to the system as the **root** user:

- a. **Install vsftpd**

Install the vsftpd package.

```
# yum install vsftpd
```

- b. **Start vsftpd**

Start the **vsftpd** daemon.

```
# chkconfig vsftpd on    service vsftpd start
```

- c. **Create sub-directory**

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

Once the FTP server has been configured, you must use the **reposync** command to download the packages to be shared. It downloads all packages from all configured software repositories.

This includes repositories for all Red Hat Network channels the system is subscribed to, and any locally configured repositories.

- a. As the **root** user, change into the **/var/ftp/pub/rhevrepo** directory.

```
# cd /var/ftp/pub/rhevrepo
```

- b. Run the **reposync** command.

```
# reposync --plugins .
```

## 6. Create local repository metadata

Use the **createrepo** command to 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;
```

## 7. Create repository configuration files

Create a **yum** configuration file, and copy it to the **/etc/yum.repos.d/** directory on client systems that you want to connect to this software repository. Ensure that the system hosting the repository is connected to the same network as the client systems where the packages are to be installed.

The configuration file can be created manually, or using a script. If using a script, then before running it you must replace **ADDRESS** in the **baseurl** with the IP address or Fully Qualified Domain Name (FQDN) of the system hosting the repository. The script must be run on this system and then distributed to the client machines. For example:

```
#!/bin/sh

REPOFILE="/etc/yum.repos.d/rhev.repo"

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;
```

## 8. Copy the repository configuration file to client system

Copy the repository configuration file to the **/etc/yum.repos.d/** directory on every system that you want to connect to this software repository. For example: Red Hat Enterprise Virtualization Manager system(s), all Red Hat Enterprise Linux virtualization hosts, and all Red Hat Enterprise Linux virtual machines.

Now that your client systems have been configured to use your local repository, you can proceed with management server, virtualization host, and virtual machine installation as documented in the Red Hat Enterprise Virtualization product documentation. Instead of installing packages from Red Hat Network, you can install them from your newly created local repository.

**NOTE**

You can also provide the software repository created here to client systems using removable media, such as a portable USB drive. To do this, first create the repository using the steps provided, and then:

1. Recursively copy the `/var/ftp/pub/rhevrepo` directory, and all its contents, to the removable media.
2. Modify the `/etc/yum.repos.d/rhev.repo` file, replacing the `baseurl` values with the path to which the removable media will be mounted on the client systems. For example `file:///media/disk/rhevrepo/`.

**NOTE**

As updated packages are released to Red Hat Network - addressing security issues, fixing bugs, and adding enhancements - you must update your local repository. To do this, repeat the procedure for synchronizing and sharing the channels. Adding the `--newest-only` parameter to the `reposync` command ensures that it only retrieves the newest version of each available package. Once the repository is updated you must ensure it is available to each of your client systems and then run `yum update` on it.

[Report a bug](#)

## 2.4.2. Installing the Red Hat Enterprise Virtualization Manager Packages

### Summary

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


### Procedure 2.3. Installing the Red Hat Enterprise 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 Enterprise Virtualization Manager:

```
# yum update
```

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

```
# yum install rhevm
```

**NOTE**

The `rhev-doc` package is installed as a dependency of the `rhev` package, and provides a local copy of the Red Hat Enterprise Virtualization documentation suite. This documentation is also used to provide context sensitive help links from the Administration and User Portals. You can run the following commands to search for translated versions of the documentation:

```
# yum search rhevm-doc
```

**Result**

You have installed the `rhev` package and dependencies.

[Report a bug](#)

**2.4.3. Configuring the Red Hat Enterprise Virtualization Manager**

After you have installed the `rhev` package and dependencies, you must configure the Red Hat Enterprise 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.

**NOTE**

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.

**Procedure 2.4. Configuring the Red Hat Enterprise Virtualization Manager****1. Packages**

The **`engine-setup`** command checks to see if it is performing an upgrade or an installation, and whether any updates are available for the packages linked to the Manager. No user input is required at this stage.

```
[ INFO ] Checking for product updates...
[ INFO ] No product updates found
```

**2. Network Configuration**

A reverse lookup is performed on the host name of the machine on which the Red Hat Enterprise Virtualization Manager is being installed. The host name is detected automatically, but you can correct this host name if it is incorrect or if you are using virtual hosts. There must be forward and reverse lookup records for the provided host name in DNS, especially if you will also install the reports server.

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

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

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

**3. Database Configuration**

You can use either a local or remote PostgreSQL database. The **`engine-setup`** command can configure your database automatically (including adding a user and a database), or it can use values that you supply.

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

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]:

#### 4. oVirt Engine Configuration

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

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

**Both** offers the greatest flexibility.

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

Engine admin password:

Confirm engine admin password:

#### 5. PKI Configuration

The Manager uses certificates to communicate securely with its hosts. You provide the organization name for the certificate. This certificate can also optionally be used to secure https communications with the Manager.

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

#### 6. Apache Configuration

By default, external SSL (HTTPS) communication with the Manager is secured with the self-signed certificate created in the PKI configuration stage to securely communicate with hosts. Another certificate may be chosen for external HTTPS connections, without affecting 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]:

The Red Hat Enterprise Virtualization Manager uses the Apache web server to present a landing page to users. The **engine-setup** command can make the landing page of the Manager the default page presented by Apache.

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]:

#### 7. System Configuration

The **engine-setup** command can create an NFS share on the Manager to use as an ISO storage domain. Hosting the ISO domain locally to the Manager simplifies keeping some elements of your environment up to date.

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

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

Local ISO domain ACL [0.0.0.0/0.0.0.0(rw)]:

Local ISO domain name [ISO\_DOMAIN]:



## 8. Websocket Proxy Server Configuration

The **engine-setup** command can optionally 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]:
```

## 9. Miscellaneous Configuration

You can use the **engine-setup** command to allow a proxy server to broker transactions from the Red Hat Access plug-in.

```
Would you like transactions from the Red Hat Access Plugin sent from the RHEV Manager to be brokered through a proxy server? (Yes, No) [No]:
```

```
[ INFO ] Stage: Setup validation
```

## 10. Configuration Preview

Check the configuration preview to confirm the values you entered before they are applied. If you choose to proceed, **engine-setup** configures the Red Hat Enterprise Virtualization Manager using those values.

```
Engine database name           : engine
Engine database secured connection : False
Engine database host           : localhost
Engine database user name      : engine
Engine database host name validation : False
Engine database port           : 5432
NFS setup                      : True
PKI organization               : Your Org
Application mode               : both
Firewall manager               : iptables
Update Firewall                : True
Configure WebSocket Proxy      : True
Host FQDN                      : Your Manager's FQDN
NFS export ACL                 : 0.0.0.0/0.0.0.0(rw)
NFS mount point                : /var/lib/exports/iso
Datacenter storage type        : nfs
Configure local Engine database : True
Set application as default page : True
Configure Apache SSL           : True
```

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

When your environment has been configured, the **engine-setup** command displays details about how to access your environment and related security details.

## 11. Clean Up and Termination

The **engine-setup** command cleans up any temporary files created during the configuration process, and outputs the location of the log file for the Red Hat Enterprise Virtualization Manager configuration process.

```
[ INFO ] Stage: Clean up
```

```
Log file is located at /var/log/ovirt-engine/setup/ovirt-engine-setup-installation-date.log
```

```
[ INFO ] Stage: Pre-termination
```

```
[ INFO ] Stage: Termination
[ INFO ] Execution of setup completed successfully
```

## Result

The Red Hat Enterprise Virtualization Manager has been configured and is running on your server. You can log in to the Administration Portal as the **admin@internal** user to continue configuring the Manager. Furthermore, the **engine-setup** command saves your answers to a file that can be used to reconfigure the Manager using the same values.

[Report a bug](#)

## 2.4.4. Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager

### Summary

You can manually configure a database server to host the database used by the Red Hat Enterprise Virtualization Manager. The database can be hosted either locally on the machine on which the Red Hat Enterprise Virtualization Manager is installed, or remotely on another machine.



### IMPORTANT

The database must be prepared prior to running the **engine-setup** command.

### Procedure 2.5. Preparing a PostgreSQL Database for use with Red Hat Enterprise Virtualization Manager

1. Run the following commands to initialize the PostgreSQL database, start the **postgresql** service and ensure this service starts on boot:

```
# service postgresql initdb
# service postgresql start
# chkconfig postgresql on
```

2. Create a user for the Red Hat Enterprise Virtualization Manager to use when it writes to and reads from the database, and a database in which to store data about the Red Hat Enterprise Virtualization environment. You must perform this step on both local and remote databases. Use the **psql** terminal as the **postgres** user.

```
# su - postgres
$ psql
postgres=# create role [user name] with login encrypted password '[password]';
postgres=# create database [database name] owner [user name] template template0
encoding 'UTF8' lc_collate 'en_US.UTF-8' lc_ctype 'en_US.UTF-8';
```

3. Run the following commands to connect to the new database and add the **plpgsql** language:

```
postgres=# \c [database name]
CREATE LANGUAGE plpgsql;
```

4. Ensure the database can be accessed remotely by enabling client authentication. Edit the **/var/lib/pgsql/data/pg\_hba.conf** file, and add the following in accordance with the location of the database:

- For local databases, add the two following lines 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
```

- For remote databases, 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:

```
host [database name] [user name] X.X.X.X/32 md5
```

5. Allow TCP/IP connections to the database. You must perform this step for remote databases. 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.

6. Restart the **postgresql** service. This step is required on both local and remote manually configured database servers.

```
# service postgresql restart
```

## Result

You have manually configured a PostgreSQL database to use with the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 2.4.5. Configuring the Manager to Use a Manually Configured Local or Remote PostgreSQL Database

### Summary

During the database configuration stage of configuring the Red Hat Enterprise Virtualization Manager using the **engine-setup** script, you can choose to use a manually configured database. You can select to use a locally or remotely installed PostgreSQL database.

### Procedure 2.6. Configuring the Manager to use a Manually Configured Local or Remote PostgreSQL Database

1. During configuration of the Red Hat Enterprise Virtualization Manager, the **engine-setup** command prompts you to decide where your database is located:

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

The steps involved in manually configuring the Red Hat Enterprise Virtualization Manager to use local or remotely hosted databases are the same. However, to use a remotely hosted database you must provide the host name of the remote database server and the port on which it is listening.

2. When prompted, enter **Manual** to manually configure the database:

Would you like Setup to automatically configure postgresql, or prefer to perform that manually? (Automatic, Manual) [Automatic]: Manual

- If you are using a remotely hosted database, supply the **engine-setup** command with the host name of your database server and the port on which it is listening:

Database host [localhost]:  
Database port [5432]:

- For both local and remotely hosted databases, you must select whether or not your database uses a secured connection. You must also enter the name of the database you configured, the user the Manager can use to access the database, and the password of that user.

Database secured connection (Yes, No) [No]:  
Database name [engine]:  
Database user [engine]:  
Database password:



#### NOTE

Using a secured connection to your database requires you to also have manually configured secured database connections.

#### Result

You have configured the Red Hat Enterprise Virtualization Manager to use a manually configured database. The **engine-setup** command continues with the rest of your environment configuration.

[Report a bug](#)


## 2.4.6. Connecting to the Administration Portal

### Summary

Access the Administration Portal using a web browser.

### Procedure 2.7. Connecting to the Administration Portal

- Open a supported 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 to open the login screen.



#### IMPORTANT

The first time that you connect to the Administration Portal, you are prompted to trust the certificate being used to secure communications between your browser and the web server.

- Enter your **User Name** and **Password**. If you are logging in for the first time, use the user name **admin** in conjunction with the administrator password that you specified during installation.

4. Select the domain against which to authenticate from the **Domain** drop-down list. If you are logging in using the internal **admin** user name, select the **internal** domain.
5. 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 list.
6. Click **Login**.

### Result

You have logged into the Administration Portal.

[Report a bug](#)

## 2.4.7. Removing the Red Hat Enterprise Virtualization Manager

### Summary

You can use the **engine-cleanup** command to remove the files associated with the Red Hat Enterprise Virtualization Manager.


### Procedure 2.8. Removing Red Hat Enterprise Virtualization Manager

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

```
# engine-cleanup
```

2. You are prompted to confirm removal of all Red Hat Enterprise Virtualization Manager components. These components include PKI keys, the locally hosted ISO domain file system layout, PKI configuration, the local NFS exports configuration, and the engine database content.

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



### 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.

3. You are given another opportunity to change your mind and cancel the removal of the Red Hat Enterprise 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
```

### Result

The configuration files of your environment have been removed according to your selections when you ran **engine-cleanup**.

```
--== SUMMARY ==--
```

```
A backup of the database is available at /var/lib/ovirt-engine/backups/engine-date-and-extra-characters.sql
```

```
Engine setup successfully cleaned up
```

```
A backup of PKI configuration and keys is available at /var/lib/ovirt-engine/backups/engine-pki-date-and-extra-characters.tar.gz
```

```
--== END OF SUMMARY ==--
```

```
[ INFO ] Generating answer file '/var/lib/ovirt-engine/setup/answers/20130827181911-cleanup.conf'
```

```
[ INFO ] Stage: Clean up
```

```
Log file is located at /var/log/ovirt-engine/setup/ovirt-engine-remove-date.log
```

```
[ INFO ] Stage: Pre-termination
```

```
[ INFO ] Stage: Termination
```

```
[ INFO ] Execution of cleanup completed successfully
```

You can now safely remove the Red Hat Enterprise Virtualization packages using the **yum** command.

```
# yum remove rhevm* vdsm-bootstrap
```

[Report a bug](#)

## 2.5. SPICE CLIENT

### 2.5.1. SPICE Features

The following SPICE features were added in the release of Red Hat Enterprise Virtualization 3.3:

#### SPICE-HTML5 support (Technology Preview), [BZ#974060](#)

Initial support for the SPICE-HTML5 console client is now offered as a technology preview. This feature allows users to connect to a SPICE console from their browser using the SPICE-HTML5 client. The requirements for enabling SPICE-HTML5 are the same as that of the noVNC console, as follows:

On the guest:

- The WebSocket proxy must be set up and running in the environment.
- The engine must be aware of the WebSocket proxy - use **engine-config** to set the **WebSocketProxy** key.

On the client:

- The client must have a browser with WebSocket and postMessage support.
- If SSL is enabled, the engine's certificate authority must be imported in the client browser.

The features of SPICE supported in each operating system depends on the version of SPICE that is packaged for that operating system.

**Table 2.1.**

Client Operating System	Wan Optimizations	Dynamic Console Resizing	SPICE Proxy Support	Full High Definition Display	Multiple Monitor Support
RHEL 5.8+	No	No	No	Yes	Yes
RHEL 6.2 - 6.4	No	No	No	Yes	Yes
RHEL 6.5 +	Yes	Yes	Yes	Yes	Yes
Windows XP (All versions)	Yes	Yes	Yes	Yes	Yes
Windows 7 (All versions)	Yes	Yes	Yes	Yes	Yes
Windows 8 (All versions)	Yes	Yes	Yes	Yes	Yes
Windows Server 2008	Yes	Yes	Yes	Yes	Yes
Windows Server 2012	Yes	Yes	Yes	Yes	Yes

[Report a bug](#)

## CHAPTER 3. THE SELF-HOSTED ENGINE

### 3.1. ABOUT THE SELF-HOSTED ENGINE

A self-hosted engine is a virtualized environment in which the engine, or Manager, runs on a virtual machine on the hosts managed by that engine. The virtual machine is created as part of the host configuration, and the engine is installed and configured in parallel to that host configuration process, referred to in these procedures as the deployment.

The virtual machine running the engine is created to be highly available. This means that if the host running the virtual machine goes into maintenance mode, or fails unexpectedly, the virtual machine will be migrated automatically to another host in the environment.

The primary benefit of the self-hosted engine is that it requires less hardware to deploy an instance of Red Hat Enterprise Virtualization as the engine runs as a virtual machine, not on physical hardware. Additionally, the engine is configured to be highly available automatically, rather than requiring a separate cluster.

The self-hosted engine currently only runs on Red Hat Enterprise Linux 6.5 or 6.6 hosts. Red Hat Enterprise Virtualization Hypervisors and older versions of Red Hat Enterprise Linux are not recommended for use with a self-hosted engine.

[Report a bug](#)

### 3.2. LIMITATIONS OF THE SELF-HOSTED ENGINE

At present there are two main limitations of the self-hosted engine configuration:

- An NFS storage domain is required for the configuration. NFS is the only supported file system for the self-hosted engine.
- The host of the self-hosted engine and all attached hosts must use Red Hat Enterprise Linux 6.5 or 6.6. Red Hat Enterprise Virtualization Hypervisors are not supported.

[Report a bug](#)

### 3.3. INSTALLING THE SELF-HOSTED ENGINE

#### Summary

Install a Red Hat Enterprise Virtualization environment that takes advantage of the self-hosted engine feature, in which the engine is installed on a virtual machine within the environment itself.

You must be subscribed to the appropriate Red Hat Network channels to install the packages. For Subscription Manager, these channels are:

- **rhel-6-server-rpms**
- **rhel-6-server-supplementary-rpms**
- **rhel-6-server-rhev-3.4-rpms**
- **jb-eap-6-for-rhel-6-server-rpms**
- **rhel-6-server-rhev-mgmt-agent-rpms**



For more information on subscribing to these channels using Subscription Manager, refer to [Section 2.3.1, “Subscribing to the Red Hat Enterprise Virtualization Manager Channels using Subscription Manager”](#).

For RHN Classic, these channels are:

- **rhel-x86\_64-server-6**
- **rhel-x86\_64-server-supplementary-6**
- **rhel-x86\_64-server-6-rhev-3.4**
- **jbappplatform-6-x86\_64-server-6-rpm**
- **rhel-x86\_64-rhev-mgmt-agent-6**

For more information on subscribing to these channels using RHN Classic, refer to [Section 2.3.2, “Subscribing to the Red Hat Enterprise Virtualization Manager Channels Using RHN Classic”](#).



### IMPORTANT

While the `ovirt-hosted-engine-setup` package is provided by the Red Hat Enterprise Virtualization Manager channel and can be installed using the standard channels for the Manager, the `vdsm` package is a dependency of the `ovirt-hosted-engine-setup` package and is provided by the Red Hat Enterprise Virt Management Agent channel, which must be enabled. This channel is **rhel-6-server-rhev-mgmt-agent-rpms** in Subscription Manager and **rhel-x86\_64-rhev-mgmt-agent-6** in RHN Classic.

All steps in this procedure are to be conducted as the **root** user.

#### Procedure 3.1. Installing the Self-Hosted Engine

1. Run the following command to ensure that the most up-to-date versions of all installed packages are in use:

```
# yum upgrade
```

2. Run the following command to install the `ovirt-hosted-engine-setup` package and dependencies:

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

#### Result

You have installed the `ovirt-hosted-engine-setup` package and are ready to configure the self-hosted engine.

[Report a bug](#)

## 3.4. CONFIGURING THE SELF-HOSTED ENGINE

### Summary

When package installation is complete, the Red Hat Enterprise Virtualization Manager must be configured. The **hosted-engine** deployment script is provided to assist with this task. The script asks

you a series of questions, and configures your environment based on your answers. When the required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.

The **hosted-engine** deployment script guides you through several distinct configuration stages. The script suggests possible configuration defaults in square brackets. Where these default values are acceptable, no additional input is required.

This procedure requires a new Red Hat Enterprise Linux 6.5 or 6.6 host with the `ovirt-hosted-engine-setup` package installed. This host is referred to as 'Host-HE1', with a fully qualified domain name (FQDN) of **Host-HE1.example.com** in this procedure.

The hosted engine, the virtual machine created during configuration of Host-HE1 to manage the environment, is referred to as 'my-engine'. You will be prompted by the **hosted-engine** deployment script to access this virtual machine multiple times to install an operating system and to configure the engine.

All steps in this procedure are to be conducted as the **root** user for the specified machine.

### Procedure 3.2. Configuring the Self-Hosted Engine

#### 1. Initiating Hosted Engine Deployment

Begin configuration of the self-hosted environment by deploying the **hosted-engine** customization script on Host\_HE1. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

```
# hosted-engine --deploy
```

#### 2. Configuring Storage

Select the version of NFS and specify the full address, using either the FQDN or IP address, and path name of the shared storage domain. Choose the storage domain and storage data center names to be used in the environment.

```
During customization use CTRL-D to abort.
Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:
Please specify the full shared storage connection path to use (example: host:/path):
storage.example.com:/hosted_engine/nfs
[ INFO ] Installing on first host
Please provide storage domain name. [hosted_storage]:
Local storage datacenter name is an internal name and currently will not be shown in
engine's admin UI.Please enter local datacenter name [hosted_datacenter]:
```

#### 3. Configuring the Network

The script detects possible network interface controllers (NICs) to use as a management bridge for the environment. It then checks your firewall configuration and offers to modify it for console (SPICE or VNC) access HostedEngine-VM. Provide a pingable gateway IP address, to be used by the **ovirt-ha-agent** to help determine a host's suitability for running HostedEngine-VM.

```
Please indicate a nic to set rhevm bridge on: (eth1, eth0) [eth1]:
iptables was detected on your computer, do you wish setup to configure it? (Yes, No)[Yes]:
Please indicate a pingable gateway IP address [X.X.X.X]:
```

#### 4. Configuring the Virtual Machine

The script creates a virtual machine to be configured as the Red Hat Enterprise Virtualization

Manager, the hosted engine referred to in this procedure as HostedEngine-VM. Specify the boot device and, if applicable, the path name of the installation media, the CPU type, the number of virtual CPUs, and the disk size. Specify a MAC address for the HostedEngine-VM, or accept a randomly generated one. The MAC address can be used to update your DHCP server prior to installing the operating system on the virtual machine. Specify memory size and console connection type for the creation of HostedEngine-VM.

Please specify the device to boot the VM from (cdrom, disk, pxe) [cdrom]:

The following CPU types are supported by this host:

- model\_Penryn: Intel Penryn Family

- model\_Conroe: Intel Conroe Family

Please specify the CPU type to be used by the VM [model\_Penryn]:

Please specify the number of virtual CPUs for the VM [Defaults to minimum requirement: 2]:

Please specify the disk size of the VM in GB [Defaults to minimum requirement: 25]:

You may specify a MAC address for the VM or accept a randomly generated default

[00:16:3e:77:b2:a4]:

Please specify the memory size of the VM in MB [Defaults to minimum requirement: 4096]:

Please specify the console type you would like to use to connect to the VM (vnc, spice) [vnc]:

## 5. Configuring the Hosted Engine

Specify the name for Host-HE1 to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user to access the Administrator Portal. Provide the FQDN for HostedEngine-VM; this procedure uses the FQDN *HostedEngine-VM.example.com*. Finally, 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.

Enter the name which will be used to identify this host inside the Administrator Portal

[hosted\_engine\_1]: Host-HE1

Enter 'admin@internal' user password that will be used for accessing the Administrator Portal:

Confirm 'admin@internal' user password:

Please provide the FQDN for the engine you would like to use. This needs to match the FQDN that you will use for the engine installation within the VM: *HostedEngine-VM.example.com*

Please provide the name of the SMTP server through which we will send notifications [localhost]:

Please provide the TCP port number of the SMTP server [25]:

Please provide the email address from which notifications will be sent [root@localhost]:

Please provide a comma-separated list of email addresses which will get notifications [root@localhost]:

## 6. Configuration Preview

Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```

Bridge interface           : eth1
Engine FQDN                : HostedEngine-VM.example.com
Bridge name                : rhevm
SSH daemon port           : 22
Firewall manager          : iptables
Gateway address           : X.X.X.X
Host name for web application : Host-HE1
Host ID                    : 1
Image size GB              : 25

```

```

Storage connection      : storage.example.com:/hosted_engine/nfs
Console type           : vnc
Memory size MB         : 4096
MAC address            : 00:16:3e:77:b2:a4
Boot type              : pxe
Number of CPUs         : 2
CPU Type               : model_Penryn

```

Please confirm installation settings (Yes, No)[No]:

## 7. Creating HostedEngine-VM

The script creates a virtual machine to be HostedEngine-VM and provides connection details. You must install an operating system on HostedEngine-VM before the **hosted-engine** script can proceed on Host-HE1.

```

[ INFO ] Generating answer file '/etc/ovirt-hosted-engine/answers.conf'
[ INFO ] Stage: Transaction setup
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Configuring libvirt
[ INFO ] Generating VDSM certificates
[ INFO ] Configuring VDSM
[ INFO ] Starting vdsmd
[ INFO ] Waiting for VDSM hardware info
[ INFO ] Creating Storage Domain
[ INFO ] Creating Storage Pool
[ INFO ] Connecting Storage Pool
[ INFO ] Verifying sanlock lockspace initialization
[ INFO ] Initializing sanlock lockspace
[ INFO ] Initializing sanlock metadata
[ INFO ] Creating VM Image
[ INFO ] Disconnecting Storage Pool
[ INFO ] Start monitoring domain
[ INFO ] Configuring VM
[ INFO ] Updating hosted-engine configuration
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up
[ INFO ] Creating VM

```

You can now connect to the VM with the following command:

```
/usr/bin/remote-viewer vnc://localhost:5900
```

Use temporary password "3042QHpx" to connect to vnc console.

Please note that in order to use remote-viewer you need to be able to run graphical applications.

This means that if you are using ssh you have to supply the -Y flag (enables trusted X11 forwarding).

Otherwise you can run the command from a terminal in your preferred desktop environment. If you cannot run graphical applications you can connect to the graphic console from another host or connect to the console using the following command:

```
virsh -c qemu+tls://Test/system console HostedEngine
```

If you need to reboot the VM you will need to start it manually using the command:

```
hosted-engine --vm-start
```

You can then set a temporary password using the command:

```
hosted-engine --add-console-password
```

The VM has been started. Install the OS and shut down or reboot it. To continue please make a selection:

- (1) Continue setup - VM installation is complete
- (2) Reboot the VM and restart installation
- (3) Abort setup

(1, 2, 3)[1]:

Using the naming convention of this procedure, connect to the virtual machine using VNC with the following command:

```
/usr/bin/remote-viewer vnc://Host-HE1.example.com:5900
```

## 8. Installing the Virtual Machine Operating System

Connect to HostedEngine-VM, the virtual machine created by the hosted-engine script, and install a Red Hat Enterprise Linux 6.5 or 6.6 operating system. Ensure the machine is rebooted once installation has completed.

## 9. Synchronizing the Host and the Virtual Machine

Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

- (1) Continue setup - VM installation is complete

Waiting for VM to shut down...

[ INFO ] Creating VM

You can now connect to the VM with the following command:

```
/usr/bin/remote-viewer vnc://localhost:5900
```

Use temporary password "3042QHpX" to connect to vnc console.

Please note that in order to use remote-viewer you need to be able to run graphical applications.

This means that if you are using ssh you have to supply the -Y flag (enables trusted X11 forwarding).

Otherwise you can run the command from a terminal in your preferred desktop environment.

If you cannot run graphical applications you can connect to the graphic console from another host or connect to the console using the following command:

```
virsh -c qemu+tls://Test/system console HostedEngine
```

If you need to reboot the VM you will need to start it manually using the command:

```
hosted-engine --vm-start
```

You can then set a temporary password using the command:

```
hosted-engine --add-console-password
```

Please install and setup the engine in the VM.

You may also be interested in subscribing to "agent" RHN/Satellite channel and installing rhvm-guest-agent-common package in the VM.

To continue make a selection from the options below:

- (1) Continue setup - engine installation is complete
- (2) Power off and restart the VM
- (3) Abort setup

## 10. Installing the Manager

Connect to HostedEngine-VM, subscribe to the appropriate Red Hat Enterprise Virtualization Manager channels, ensure that the most up-to-date versions of all installed packages are in use, and install the rhvm packages.

```
# yum upgrade
```

```
# yum install rhevm
```

### 11. Configuring the Manager

Configure the engine on HostedEngine-VM:

```
# engine-setup
```

### 12. Synchronizing the Host and the Manager

Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

```
(1) Continue setup - engine installation is complete
```

```
[ INFO ] Engine replied: DB Up!Welcome to Health Status!  
[ INFO ] Waiting for the host to become operational in the engine. This may take several  
minutes...  
[ INFO ] Still waiting for VDSM host to become operational...  
[ INFO ] The VDSM Host is now operational  
Please shutdown the VM allowing the system to launch it as a monitored service.  
The system will wait until the VM is down.
```

### 13. Shutting Down HostedEngine-VM

Shutdown HostedEngine-VM.

```
# shutdown now
```

### 14. Setup Confirmation

Return to Host-HE1 to confirm it has detected that HostedEngine-VM is down.

```
[ INFO ] Enabling and starting HA services  
Hosted Engine successfully set up  
[ INFO ] Stage: Clean up  
[ INFO ] Stage: Pre-termination  
[ INFO ] Stage: Termination
```

## Result

When the **hosted-engine** deployment script completes successfully, the Red Hat Enterprise Virtualization Manager is configured and running on your server. In contrast to a bare-metal Manager installation, the hosted engine Manager has already configured the data center, cluster, host (Host-HE1), storage domain, and virtual machine of the hosted engine (HostedEngine-VM). You can log in as the **admin@internal** user to continue configuring the Manager and add further resources.

Link your Red Hat Enterprise Virtualization Manager to a directory server so you can add additional users to the environment. Red Hat Enterprise Virtualization supports directory services from Red Hat Directory Services (RHDS), IdM, and Active Directory. Add a directory server to your environment using the **engine-manage-domains** command.

The **ovirt-host-engine-setup** script also saves the answers you gave during configuration to a file, to help with disaster recovery. If a destination is not specified using the **--generate-answer=<file>** argument, the answer file is generated at **/etc/ovirt-hosted-engine/answers.conf**.

[Report a bug](#)

## 3.5. INSTALLING ADDITIONAL HOSTS TO A SELF-HOSTED ENVIRONMENT

### Summary

Adding additional nodes to a self-hosted environment is very similar to deploying the original host, though heavily truncated as the script detects the environment.

As with the original host, additional hosts require Red Hat Enterprise Linux 6.5 or 6.6 with subscriptions to the appropriate Red Hat Enterprise Virtualization channels.

All steps in this procedure are to be conducted as the **root** user.

### Procedure 3.3. Adding the host

1. Install the `ovirt-hosted-engine-setup` package.

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

2. Configure the host with the deployment command.

```
# hosted-engine --deploy
```

3. **Configuring Storage**

Specify the storage type and the full address, using either the Fully Qualified Domain Name (FQDN) or IP address, and path name of the shared storage domain used in the self-hosted environment.

```
Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:
Please specify the full shared storage connection path to use (example: host:/path):
storage.example.com:/hosted_engine/nfs
```

4. **Detecting the Self-Hosted Engine**

The **hosted-engine** script detects that the shared storage is being used and asks if this is an additional host setup. You are then prompted for the host ID, which must be an integer not already assigned to an additional host in the environment.

```
The specified storage location already contains a data domain. Is this an additional host
setup (Yes, No)[Yes]?
[ INFO ] Installing on additional host
Please specify the Host ID [Must be integer, default: 2]:
```

5. **Configuring the System**

The **hosted-engine** script uses the answer file generated by the original hosted-engine setup. To achieve this, the script requires the FQDN or IP address and the password of the **root** user of that host so as to access and secure-copy the answer file to the additional host.

```
[WARNING] A configuration file must be supplied to deploy Hosted Engine on an additional
host.
The answer file may be fetched from the first host using scp.
If you do not want to download it automatically you can abort the setup answering no to the
following question.
Do you want to scp the answer file from the first host? (Yes, No)[Yes]:
```

```
Please provide the FQDN or IP of the first host:
Enter 'root' user password for host Host-HE1.example.com:
[ INFO ] Answer file successfully downloaded
```

## 6. Configuring the Hosted Engine

Specify the name for the additional host to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user.

```
Enter the name which will be used to identify this host inside the Administrator Portal
[hosted_engine_2]:
Enter 'admin@internal' user password that will be used for accessing the Administrator
Portal:
Confirm 'admin@internal' user password:
```

## 7. Configuration Preview

Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface           : eth1
Engine FQDN                : HostedEngine-VM.example.com
Bridge name                : rhevm
SSH daemon port            : 22
Firewall manager           : iptables
Gateway address            : X.X.X.X
Host name for web application : hosted_engine_2
Host ID                    : 2
Image size GB              : 25
Storage connection         : storage.example.com:/hosted_engine/nfs
Console type               : vnc
Memory size MB             : 4096
MAC address                : 00:16:3e:05:95:50
Boot type                  : disk
Number of CPUs             : 2
CPU Type                   : model_Penryn

Please confirm installation settings (Yes, No)[Yes]:
```

### Result

After confirmation, the script completes installation of the host and adds it to the environment.

[Report a bug](#)

## 3.6. MAINTAINING THE SELF-HOSTED ENGINE

The maintenance modes enable you to start, stop, and modify the engine virtual machine without interference from the high-availability agents, and to restart and modify the hosts in the environment without interfering with the engine.

There are three maintenance modes that can be enforced:

- **global** - All high-availability agents in the cluster are disabled from monitoring the state of the engine virtual machine. The **global** maintenance mode must be applied for any setup or upgrade operations that require the engine to be stopped. Examples of this include upgrading



to a later version of Red Hat Enterprise Virtualization, and installation of the `rhev-dwh` and `rhev-reports` packages necessary for the **Reports Portal**.

- **local** - The high-availability agent on the host issuing the command is disabled from monitoring the state of the engine virtual machine. The host is exempt from hosting the engine virtual machine while in **local** maintenance mode; if hosting the engine virtual machine when placed into this mode, the engine will be migrated to another host, provided there is a suitable contender. The **local** maintenance mode is recommended when applying system changes or updates to the host.
- **none** - Disables maintenance mode, ensuring that the high-availability agents are operating.

The syntax for maintenance mode is:

```
# hosted-engine --set-maintenance --mode=mode
```

This command is to be conducted as the **root** user.

[Report a bug](#)

## 3.7. UPGRADING THE SELF-HOSTED ENGINE

### Summary

Upgrade your Red Hat Enterprise Virtualization hosted-engine environment from version 3.3 to 3.4.

This procedure upgrades two hosts, referred to in this procedure as Host A and Host B, and a Manager virtual machine. For the purposes of this procedure, Host B is hosting the Manager virtual machine.

It is recommended that all hosts in the environment be upgraded at the same time, before the Manager virtual machine is upgraded and the **Compatibility Version** of the cluster is updated to **3.4**. This avoids any version 3.3 hosts from going into a **Non Operational** state.


All commands in this procedure are as the **root** user.

### Procedure 3.4. Upgrading the Self-Hosted Engine

1. Log into either host and set the maintenance mode to **global** to disable the high-availability agents.

```
# hosted-engine --set-maintenance --mode=global
```

2. Access the Red Hat Enterprise Virtualization Manager Administration Portal. Select Host A and put it into maintenance mode by clicking the **Maintenance** button.



#### IMPORTANT

The host that you put into maintenance mode and upgrade must not be the host currently hosting the Manager virtual machine.

3. Log into and update Host A.

```
# yum update
```

- Restart VDSM on Host A.

```
# service vdsmd restart
```

- Restart **ovirt-ha-broker** and **ovirt-ha-agent** on Host A.

```
# service ovirt-ha-broker restart
```

```
# service ovirt-ha-agent restart
```

- Log into either host and turn off the hosted-engine maintenance mode so that the Manager virtual machine can migrate to the other host.

```
# hosted-engine --set-maintenance --mode=none
```

- Access the Red Hat Enterprise Virtualization Manager Administration Portal. Select Host A and activate it by clicking the **Activate** button.

- Log into Host B and set the maintenance mode to **global** to disable the high-availability agents.

```
# hosted-engine --set-maintenance --mode=global
```

- Update Host B.

```
# yum update
```

- Restart VDSM on Host B.

```
# service vdsmd restart
```

- Restart **ovirt-ha-broker** and **ovirt-ha-agent** on Host B.

```
# service ovirt-ha-broker restart
```

```
# service ovirt-ha-agent restart
```

- Turn off the hosted-engine maintenance mode on Host B.

```
# hosted-engine --set-maintenance --mode=none
```

- Access the Red Hat Enterprise Virtualization Manager Administration Portal. Select Host B and activate it by clicking the **Activate** button.

- Log into the Manager virtual machine and update the engine as per the instructions in [Section 5.2.4, "Upgrading to Red Hat Enterprise Virtualization Manager 3.4"](#).

- Access the Red Hat Enterprise Virtualization Manager Administration Portal.

- Select the **Default** cluster and click **Edit** to open the **Edit Cluster** window.

- Use the **Compatibility Version** drop-down menu to select **3.4**. Click **OK** to save the change and close the window.

## Result

You have upgraded both the hosts and the Manager in your hosted-engine setup to Red Hat Enterprise Virtualization 3.4.

[Report a bug](#)

## 3.8. UPGRADING ADDITIONAL HOSTS IN A SELF-HOSTED ENVIRONMENT

### Summary

It is recommended that all hosts in your self-hosted environment are upgraded at the same time. This prevents version 3.3 hosts from going into a **Non Operational** state. If this is not practical in your environment, follow this procedure to upgrade any additional hosts.

Ensure the host is not hosting the Manager virtual machine before beginning the procedure.

All commands in this procedure are as the **root** user.

### Procedure 3.5. Upgrading Additional Hosts

1. Log into the host and set the maintenance mode to **local**.

```
# hosted-engine --set-maintenance --mode=local
```

2. Access the Red Hat Enterprise Virtualization Manager Administration Portal. Select the host and put it into maintenance mode by clicking the **Maintenance** button.

3. Log into and update the host.

```
# yum update
```

4. Restart VDSM on the host.

```
# service vdsmd restart
```

5. Restart **ovirt-ha-broker** and **ovirt-ha-agent** on the host.

```
# service ovirt-ha-broker restart
```

```
# service ovirt-ha-agent restart
```

6. Turn off the hosted-engine maintenance mode on the host.

```
# hosted-engine --set-maintenance --mode=none
```

7. Access the Red Hat Enterprise Virtualization Manager Administration Portal. Select the host and activate it by clicking the **Activate** button.

## Result

You have updated an additional host in your self-hosted environment to Red Hat Enterprise Virtualization 3.4.

[Report a bug](#)

## 3.9. BACKING UP AND RESTORING A SELF-HOSTED ENVIRONMENT

This section explains how to back up a self-hosted engine environment and restore it on a freshly installed host. The supported backup method uses the **engine-backup** tool, and only allows you to back up the Red Hat Enterprise Virtualization Manager virtual machine but not the host that contains the Manager virtual machine.

Backing up and restoring a self-hosted engine environment involves the following key actions:

1. Back up the original Red Hat Enterprise Virtualization Manager configuration settings and database content.
2. Create a freshly installed Red Hat Enterprise Linux host and run the hosted-engine deployment script.
3. Restore the Red Hat Enterprise Virtualization Manager configuration settings and database content in the new Manager virtual machine.
4. Removing hosted-engine hosts in a **Non Operational** state and re-installing them into the restored self-hosted engine environment.

### Prerequisites

- To restore a self-hosted engine environment, you must prepare a freshly installed Red Hat Enterprise Linux system on a physical host.
- The operating system version of the new host and Manager must be the same as that of the original host and Manager.
- You must have entitlements to subscribe your new environment. For a list of the required repositories, see [Subscribing to the Required Entitlements](#).
- The fully qualified domain name of the new Manager must be the same fully qualified domain name as that of the original Manager. Forward and reverse lookup records must both be set in DNS.
- The new Manager database must have the same database user name as the original Manager database.

[Report a bug](#)

### 3.9.1. Backing up the Self-Hosted Engine Manager Virtual Machine

It is recommended that you back up your self-hosted engine environment regularly. The supported backup method uses the **engine-backup** tool and can be performed without interrupting the **ovirt-engine** service. The **engine-backup** tool only allows you to back up the Red Hat Enterprise Virtualization Manager virtual machine, but not the host that contains the Manager virtual machine.

#### Procedure 3.6. Backing up the Original Red Hat Enterprise Virtualization Manager

### 1. Preparing the Failover Host

A failover host, one of the hosted-engine hosts in the environment, must be placed into maintenance mode so that it has no virtual load at the time of the backup. This host can then later be used to deploy the restored self-hosted engine environment. Any of the hosted-engine hosts can be used as the failover host for this backup scenario, however the restore process is more straightforward if **Host 1** is used. The default name for the **Host 1** host is **hosted\_engine\_1**; this was set when the hosted-engine deployment script was initially run.

- a. Log in to one of the hosted-engine hosts.
- b. Confirm that the **hosted\_engine\_1** host is **Host 1**:

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

- c. Log in to the Administration Portal.
- d. Select the **Hosts** tab.
- e. Select the **hosted\_engine\_1** host in the results list, and click **Maintenance**.
- f. Click **Ok**.

### 2. Disabling the High-Availability Agents

Disable the high-availability agents on the hosted-engine hosts to prevent migration of the Red Hat Enterprise Virtualization Manager virtual machine during the backup process. Connect to any of the hosted-engine hosts and place the high-availability agents on all hosts into global maintenance mode.

```
# hosted-engine --set-maintenance --mode=global
```

### 3. Creating a Backup of the Manager

On the Manager virtual machine, back up the configuration settings and database content, replacing *[EngineBackupFile]* with the file name for the backup file, and *[LogFILE]* with the file name for the backup log.

```
# engine-backup --mode=backup --file=[EngineBackupFile] --log=[LogFILE]
```

### 4. Copying the Backup Files to an External Server

Securely copy the backup 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. This step is not mandatory, but the backup files must be accessible to restore the configuration settings and database content.

```
# scp -p [EngineBackupFiles] [Storage.example.com:/backup/EngineBackupFiles]
```

### 5. Enabling the High-Availability Agents

Connect to any of the hosted-engine hosts and turn off global maintenance mode. This enables the high-availability agents.

```
# hosted-engine --set-maintenance --none
```

### 6. Activating the Failover Host

Bring the **hosted\_engine\_1** host out of maintenance mode.


- a. Log in to the Administration Portal.
- b. Select the **Hosts** tab.
- c. Select **hosted\_engine\_1** from the results list.
- d. Click **Activate**.

You have backed up the configuration settings and database content of the Red Hat Enterprise Virtualization Manager virtual machine.

[Report a bug](#)

### 3.9.2. Creating a New Self-Hosted Engine Environment to be Used as the Restored Environment

You can restore a self-hosted engine on hardware that was used in the backed-up environment. However, you must use the failover host for the restored deployment. The failover host, **Host 1**, used in [Section 3.9.1, “Backing up the Self-Hosted Engine Manager Virtual Machine”](#) uses the default hostname of **hosted\_engine\_1** which is also used in this procedure. Due to the nature of the restore process for the self-hosted engine, before the final synchronization of the restored engine can take place, this failover host will need to be removed, and this can only be achieved if the host had no virtual load when the backup was taken. You can also restore the backup on a separate hardware which was not used in the backed up environment and this is not a concern.



#### IMPORTANT

This procedure assumes that you have a freshly installed Red Hat Enterprise Linux system on a physical host, have subscribed the host to the required entitlements, and installed the `ovirt-hosted-engine-setup` package. See [Section 3.3, “Installing the Self-Hosted Engine”](#) for more information.

#### Procedure 3.7. Creating a New Self-Hosted Environment to be Used as the Restored Environment

##### 1. Updating DNS

Update your DNS so that the fully qualified domain name of the Red Hat Enterprise Virtualization environment correlates to the IP address of the new Manager. In this procedure, fully qualified domain name was set as *Manager.example.com*. The fully qualified domain name provided for the engine must be identical to that given in the engine setup of the original engine that was backed up.

##### 2. Initiating Hosted Engine Deployment

On the newly installed Red Hat Enterprise Linux host, run the **hosted-engine** deployment script. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

```
# hosted-engine --deploy
```

If running the **hosted-engine** deployment script over a network, it is recommended to use the **screen** window manager to avoid losing the session in case of network or terminal disruption. Install the `screen` package first if not installed.

```
# screen hosted-engine --deploy
```

### 3. Configuring Storage

Select the type of storage to use.

During customization use CTRL-D to abort.

Please specify the storage you would like to use (iscsi, nfs3, nfs4)[nfs3]:

- For NFS storage types, specify the full address, using either the fully qualified domain name or IP address, and path name of the shared storage domain.

Please specify the full shared storage connection path to use (example: host:/path):  
*storage.example.com:/hosted\_engine/nfs*

- For iSCSI, specify the iSCSI portal IP address, port, user name and password, and select a target name from the auto-detected list:

Please specify the iSCSI portal IP address:

Please specify the iSCSI portal port [3260]:

Please specify the iSCSI portal user:

Please specify the iSCSI portal password:

Please specify the target name (*auto-detected values*) [*default*]:

Choose the storage domain and storage data center names to be used in the environment.

[ INFO ] Installing on first host

Please provide storage domain name. [hosted\_storage]:

Local storage datacenter name is an internal name and currently will not be shown in engine's admin UI. Please enter local datacenter name [hosted\_datacenter]:

### 4. Configuring the Network

The script detects possible network interface controllers (NICs) to use as a management bridge for the environment. It then checks your firewall configuration and offers to modify it for console (SPICE or VNC) access the Manager virtual machine. Provide a pingable gateway IP address, to be used by the **ovirt-ha-agent**, to help determine a host's suitability for running a Manager virtual machine.

Please indicate a nic to set rhevm bridge on: (eth1, eth0) [eth1]:

iptables was detected on your computer, do you wish setup to configure it? (Yes, No)[Yes]:

Please indicate a pingable gateway IP address [X.X.X.X]:

### 5. Configuring the New Manager Virtual Machine

The script creates a virtual machine to be configured as the new Manager virtual machine. Specify the boot device and, if applicable, the path name of the installation media, the CPU type, the number of virtual CPUs, and the disk size. Specify a MAC address for the Manager virtual machine, or accept a randomly generated one. The MAC address can be used to update your DHCP server prior to installing the operating system on the Manager virtual machine. Specify memory size and console connection type for the creation of Manager virtual machine.

Please specify the device to boot the VM from (cdrom, disk, pxe) [cdrom]:

The following CPU types are supported by this host:

- model\_Penryn: Intel Penryn Family

- model\_Conroe: Intel Conroe Family

Please specify the CPU type to be used by the VM [model\_Penryn]:

Please specify the number of virtual CPUs for the VM [Defaults to minimum requirement: 2]:

Please specify the disk size of the VM in GB [Defaults to minimum requirement: 25]:  
 You may specify a MAC address for the VM or accept a randomly generated default [00:16:3e:77:b2:a4]:  
 Please specify the memory size of the VM in MB [Defaults to minimum requirement: 4096]:  
 Please specify the console type you want to use to connect to the VM (vnc, spice) [vnc]:

## 6. Identifying the Name of the Host

A unique name must be provided for the name of the host, to ensure that it does not conflict with other resources that will be present when the engine has been restored from the backup. The name **hosted\_engine\_1** can be used in this procedure because this host was placed into maintenance mode before the environment was backed up, enabling removal of this host between the restoring of the engine and the final synchronization of the host and the engine.

Enter the name which will be used to identify this host inside the Administrator Portal [hosted\_engine\_1]:

## 7. Configuring the Hosted Engine

Specify a name for the self-hosted engine environment, and the password for the **admin@internal** user to access the Administrator Portal. Provide the fully qualified domain name for the new Manager virtual machine. This procedure uses the fully qualified domain name *Manager.example.com*. 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.



### IMPORTANT

The fully qualified domain name provided for the engine (*Manager.example.com*) must be the same fully qualified domain name provided when original Manager was initially set up.

Enter 'admin@internal' user password that will be used for accessing the Administrator Portal:  
 Confirm 'admin@internal' user password:  
 Please provide the FQDN for the engine you want to use. This needs to match the FQDN that you will use for the engine installation within the VM: *Manager.example.com*  
 Please provide the name of the SMTP server through which we will send notifications [localhost]:  
 Please provide the TCP port number of the SMTP server [25]:  
 Please provide the email address from which notifications will be sent [root@localhost]:  
 Please provide a comma-separated list of email addresses which will get notifications [root@localhost]:

## 8. Configuration Preview

Before proceeding, the **hosted-engine** deployment script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface           : eth1
Engine FQDN                : Manager.example.com
Bridge name                : rhevm
SSH daemon port           : 22
Firewall manager          : iptables
Gateway address           : X.X.X.X
Host name for web application : hosted_engine_1
```



```

Host ID                : 1
Image size GB          : 25
Storage connection    : storage.example.com:/hosted_engine/nfs
Console type          : vnc
Memory size MB         : 4096
MAC address            : 00:16:3e:77:b2:a4
Boot type              : pxe
Number of CPUs         : 2
CPU Type               : model_Penryn

```

Please confirm installation settings (Yes, No)[No]:

## 9. Creating the New Manager Virtual Machine

The script creates the virtual machine to be configured as the Manager virtual machine and provides connection details. You must install an operating system on it before the **hosted-engine** deployment script can proceed on Hosted Engine configuration.

```

[ INFO ] Generating answer file '/etc/ovirt-hosted-engine/answers.conf'
[ INFO ] Stage: Transaction setup
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Configuring libvirt
[ INFO ] Generating VDSM certificates
[ INFO ] Configuring VDSM
[ INFO ] Starting vdsmd
[ INFO ] Waiting for VDSM hardware info
[ INFO ] Creating Storage Domain
[ INFO ] Creating Storage Pool
[ INFO ] Connecting Storage Pool
[ INFO ] Verifying sanlock lockspace initialization
[ INFO ] Initializing sanlock lockspace
[ INFO ] Initializing sanlock metadata
[ INFO ] Creating VM Image
[ INFO ] Disconnecting Storage Pool
[ INFO ] Start monitoring domain
[ INFO ] Configuring VM
[ INFO ] Updating hosted-engine configuration
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up
[ INFO ] Creating VM
You can now connect to the VM with the following command:
/usr/bin/remote-viewer vnc://localhost:5900
Use temporary password "5379skAb" to connect to vnc console.
Please note that in order to use remote-viewer you need to be able to run graphical
applications.
This means that if you are using ssh you have to supply the -Y flag (enables trusted X11
forwarding).
Otherwise you can run the command from a terminal in your preferred desktop environment.
If you cannot run graphical applications you can connect to the graphic console from another
host or connect to the console using the following command:
virsh -c qemu+tls://Test/system console HostedEngine
If you need to reboot the VM you will need to start it manually using the command:
hosted-engine --vm-start
You can then set a temporary password using the command:
hosted-engine --add-console-password

```

The VM has been started. Install the OS and shut down or reboot it. To continue please make a selection:

- (1) Continue setup - VM installation is complete
- (2) Reboot the VM and restart installation
- (3) Abort setup

(1, 2, 3)[1]:

Using the naming convention of this procedure, connect to the virtual machine using VNC with the following command:

```
/usr/bin/remote-viewer vnc://hosted_engine_1.example.com:5900
```

## 10. Installing the Virtual Machine Operating System

Connect to Manager virtual machine and install a Red Hat Enterprise Linux 6.5 or 6.6 operating system.

## 11. Synchronizing the Host and the Manager

Return to the host and continue the **hosted-engine** deployment script by selecting option 1:

- (1) Continue setup - VM installation is complete

Waiting for VM to shut down...

[ INFO ] Creating VM

You can now connect to the VM with the following command:

```
/usr/bin/remote-viewer vnc://localhost:5900
```

Use temporary password "5379skAb" to connect to vnc console.

Please note that in order to use remote-viewer you need to be able to run graphical applications.

This means that if you are using ssh you have to supply the -Y flag (enables trusted X11 forwarding).

Otherwise you can run the command from a terminal in your preferred desktop environment.

If you cannot run graphical applications you can connect to the graphic console from another host or connect to the console using the following command:

```
virsh -c qemu+tls://Test/system console HostedEngine
```

If you need to reboot the VM you will need to start it manually using the command:

```
hosted-engine --vm-start
```

You can then set a temporary password using the command:

```
hosted-engine --add-console-password
```

Please install and setup the engine in the VM.

You may also be interested in subscribing to "agent" RHN/Satellite channel and installing `rhev-guest-agent-common` package in the VM.

To continue make a selection from the options below:

- (1) Continue setup - engine installation is complete
- (2) Power off and restart the VM
- (3) Abort setup

## 12. Installing the Manager

Connect to new Manager virtual machine, ensure the latest versions of all installed packages are in use, and install the `rhev` packages.

```
# yum upgrade
```

```
# yum install rhvm
```

### 13. Install Reports and the Data Warehouse

If you are also restoring Reports and the Data Warehouse, install the `rhvm-reports-setup` and `rhvm-dwh-setup` packages.

```
# yum install rhvm-reports-setup rhvm-dwh-setup
```

After the packages have completed installation, you will be able to continue with restoring the self-hosted engine Manager.

[Report a bug](#)

### 3.9.3. Restoring the Self-Hosted Engine Manager

The following procedure outlines how to restore the configuration settings and database content for a backed-up self-hosted engine Manager virtual machine.

#### Procedure 3.8. Restoring the Self-Hosted Engine Manager

1. Manually create an empty database to which the database content in the backup can be restored. The following steps must be performed on the machine where the database is to be hosted.
  1. If the database is to be hosted on a machine other than the Manager virtual machine, install the `postgresql-server` package. This step is not required if the database is to be hosted on the Manager virtual machine because this package is included with the `rhvm` package.

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

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

```
# service postgresql initdb
# service postgresql start
# chkconfig postgresql on
```

3. Enter the `postgresql` command line:

```
# su postgres
$ psql
```

4. Create the **engine** user:

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

If you are also restoring the Reports and Data Warehouse, create the **ovirt\_engine\_reports** and **ovirt\_engine\_history** users on the relevant host:

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

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

- 
- 
- 
- 
5. Create the new database:

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

If you are also restoring the Reports and Data Warehouse, create the databases on the relevant host:

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

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

- 
- 
- 
- 
- 
6. Exit the postgresql command line and log out of the postgres user:

```
postgres=# \q
$ exit
```

- 
- 
- 
- 
- 
- 
7. Edit the `/var/lib/pgsql/data/pg_hba.conf` file as follows:

- For each local database, replace the existing directives in the section starting with **local** at the bottom of the file with the following directives:

```
host database_name user_name 0.0.0.0/0 md5
host database_name user_name ::0/0 md5
```

- For each remote database:

- 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:

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

- 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.

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

```
# iptables -I INPUT 5 -p tcp -s Manager_IP_Address --dport 5432 -j ACCEPT
# service iptables save
```

- 
- 
- 
- 
- 
- 
- 
8. Restart the **postgresql** service:

```
# service postgresql restart
```

## 2. Copying the Backup Files to the New Manager

Securely copy the backup files to the new Manager virtual machine. This example copies the files from a network storage server to which the files were copied as in [Section 3.9.1, “Backing up the Self-Hosted Engine Manager Virtual Machine”](#). In this example, `[Storage.example.com]` is the fully qualified domain name of the storage server, `[/backup/EngineBackupFiles]` is the designated file path for the backup files on the storage server, and `[/backup/]` is the file path for the files to which the files will be copied on the new Manager.

```
# scp -p [Storage.example.com:/backup/EngineBackupFiles] [/backup/]
```

- Restore a complete backup or a database-only backup with the **`--change-db-credentials`** parameter to pass the credentials of the new database. The `database_location` for a database local to the Manager is **localhost**.



### NOTE

The following examples use a **`--*password`** option for each database without specifying a password, which will prompt for a password for each database. Passwords can be supplied for these options in the command itself, however this is not recommended as the password will then be stored in the shell history. Alternatively, **`--*passfile=password_file`** options can be used for each database to securely pass the passwords to the **engine-backup** tool without the need for interactive prompts.

- Restore a complete backup:

```
# engine-backup --mode=restore --file=file_name --log=log_file_name --change-db-credentials --db-host=database_location --db-name=database_name --db-user=engine -db-password
```

If Reports and Data Warehouse are also being restored as part of the complete backup, include the revised credentials for the two additional databases:

```
engine-backup --mode=restore --file=file_name --log=log_file_name --change-db-credentials --db-host=database_location --db-name=database_name --db-user=engine -db-password --change-reports-db-credentials --reports-db-host=database_location --reports-db-name=database_name --reports-db-user=ovirt_engine_reports --reports-db-password --change-dwh-db-credentials --dwh-db-host=database_location --dwh-db-name=database_name --dwh-db-user=ovirt_engine_history --dwh-db-password
```

- Restore a database-only backup by first restoring the configuration files backup and then restoring the database backup:

```
# engine-backup --mode=restore --scope=files --file=file_name --log=log_file_name
```

```
# engine-backup --mode=restore --scope=db --file=file_name --log=file_name --change-db-credentials --db-host=database_location --db-name=database_name --db-user=engine --db-password
```

The example above restores a backup of the Manager database.

```
# engine-backup --mode=restore --scope=reportsdb --file=file_name --log=file_name --
change-reports-db-credentials --reports-db-host=database_location --reports-db-
name=database_name --reports-db-user=ovirt_engine_reports --reports-db-password
```

The example above restores a backup of the Reports database.

```
# engine-backup --mode=restore --scope=dwhdb --file=file_name --log=file_name --
change-dwh-db-credentials --dwh-db-host=database_location --dwh-db-
name=database_name --dwh-db-user=ovirt_engine_history --dwh-db-password
```

The example above restores a backup of the Data Warehouse database.

If successful, the following output displays:

```
You should now run engine-setup.
Done.
```

#### 4. Configuring the Manager

Configure the restored Manager virtual machine. This process identifies the existing configuration settings and database content. Confirm the settings. Upon completion, the setup provides an SSH fingerprint and an internal Certificate Authority hash.

```
# engine-setup

[ INFO ] Stage: Initializing
[ INFO ] Stage: Environment setup
Configuration files: ['/etc/ovirt-engine-setup.conf.d/10-packaging.conf', '/etc/ovirt-engine-
setup.conf.d/20-setup-ovirt-post.conf']
Log file: /var/log/ovirt-engine/setup/ovirt-engine-setup-20140304075238.log
Version: otopi-1.1.2 (otopi-1.1.2-1.el6ev)
[ INFO ] Stage: Environment packages setup
[ INFO ] Yum Downloading: rhel-65-zstream/primary_db 2.8 M(70%)
[ INFO ] Stage: Programs detection
[ INFO ] Stage: Environment setup
[ INFO ] Stage: Environment customization

==== PACKAGES ====

[ INFO ] Checking for product updates...
[ INFO ] No product updates found

==== NETWORK CONFIGURATION ====

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]:
[ INFO ] iptables will be configured as firewall manager.

==== DATABASE CONFIGURATION ====

==== OVIRT ENGINE CONFIGURATION ====

Skipping storing options as database already prepared
```

```
==== PKI CONFIGURATION ====
```

```
PKI is already configured
```

```
==== APACHE CONFIGURATION ====
```

```
==== SYSTEM CONFIGURATION ====
```

```
==== END OF CONFIGURATION ====
```

```
[ INFO ] Stage: Setup validation
```

```
[WARNING] Less than 16384MB of memory is available
```

```
[ INFO ] Cleaning stale zombie tasks
```

```
==== CONFIGURATION PREVIEW ====
```

```
Database name           : engine
Database secured connection : False
Database host           : X.X.X.X
Database user name      : engine
Database host name validation : False
Database port           : 5432
NFS setup                : True
Firewall manager        : iptables
Update Firewall         : True
Configure WebSocket Proxy : True
Host FQDN                : Manager.example.com
NFS mount point         : /var/lib/exports/iso
Set application as default page : True
Configure Apache SSL    : True
```

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

## 5. Removing the Host from the Restored Environment

If the deployment of the restored self-hosted engine is on new hardware that has a unique name not present in the backed-up engine, skip this step. This step is only applicable to deployments occurring on the failover host, **hosted\_engine\_1**. Because this host was present in the environment at time the backup was created, it maintains a presence in the restored engine and must first be removed from the environment before final synchronization can take place.

- a. Log in to the Administration Portal.
- b. Click the **Hosts** tab. The failover host, **hosted\_engine\_1**, will be in maintenance mode and without a virtual load, as this was how it was prepared for the backup.
- c. Click **Remove**.
- d. Click **Ok**.

## 6. Synchronizing the Host and the Manager

Return to the host and continue the **hosted-engine** deployment script by selecting option 1:

- (1) Continue setup - engine installation is complete

```
[ INFO ] Engine replied: DB Up!Welcome to Health Status!
[ INFO ] Waiting for the host to become operational in the engine. This may take several
minutes...
[ INFO ] Still waiting for VDSM host to become operational...
```

At this point, **hosted\_engine\_1** will become visible in the Administration Portal with **Installing** and **Initializing** states before entering a **Non Operational** state. The host will continue to wait for VDSM host to become operational until it eventually times out. This happens because another host in the environment maintains the Storage Pool Manager (SPM) role and **hosted\_engine\_1** cannot interact with the storage domain because the host with SPM is in a **Non Responsive** state. When this process times out, you are prompted to shut down the virtual machine to complete the deployment. When deployment is complete, the host can be manually placed into maintenance mode and activated through the Administration Portal.

```
[ INFO ] Still waiting for VDSM host to become operational...
[ ERROR ] Timed out while waiting for host to start. Please check the logs.
[ ERROR ] Unable to add hosted_engine_2 to the manager
Please shutdown the VM allowing the system to launch it as a monitored service.
The system will wait until the VM is down.
```

## 7. Shutting Down the Manager

Shutdown the new Manager virtual machine.

```
# shutdown -h now
```

## 8. Setup Confirmation

Return to the host to confirm it has detected that the Manager virtual machine is down.

```
[ INFO ] Enabling and starting HA services
Hosted Engine successfully set up
[ INFO ] Stage: Clean up
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination
```

## 9. Activating the Host

- a. Log in to the Administration Portal.
- b. Click the **Hosts** tab.
- c. Select **hosted\_engine\_1** and click the **Maintenance** button. The host may take several maintenance before it enters maintenance mode.
- d. Click the **Activate** button.

Once active, **hosted\_engine\_1** immediately contends for SPM, and the storage domain and data center become active.

## 10. Migrating Virtual Machines to the Active Host

Migrate virtual machines to the active host by manually fencing the **Non Responsive** hosts. In the Administration Portal, right-click the hosts and select **Confirm 'Host has been Rebooted'**.

Any virtual machines that were running on that host at the time of the backup will now be removed from that host, and move from an **Unknown** state to a **Down** state. These virtual



machines can now be run on **hosted\_engine\_1**. The host that was fenced can now be forcefully removed using the REST API.

The environment has now been restored to a point where **hosted\_engine\_1** is active and is able to run virtual machines in the restored environment. The remaining hosted-engine hosts in **Non Operational** state can now be removed and re-installed into the environment.

[Report a bug](#)

### 3.9.4. Removing Non-Operational Hosts from a Restored Self-Hosted Engine Environment

Once a host has been fenced in the Administration Portal, it can be forcefully removed with a REST API request. This procedure will use **cURL**, a command line interface for sending requests to HTTP servers. Most Linux distributions include **cURL**. This procedure will connect to the Manager virtual machine to perform the relevant requests.

#### 1. Fencing the Non-Operational Host

In the Administration Portal, right-click the hosts and select **Confirm 'Host has been Rebooted'**.

Any virtual machines that were running on that host at the time of the backup will now be removed from that host, and move from an **Unknown** state to a **Down** state. The host that was fenced can now be forcefully removed using the REST API.

#### 2. Retrieving the Manager Certificate Authority

Connect to the Manager virtual machine and use the command line to perform the following requests with **cURL**.

Use a **GET** request to retrieve the Manager Certificate Authority (CA) certificate for use in all future API requests. In the following example, the **--output** option is used to designate the file *hosted-engine.ca* as the output for the Manager CA certificate. The **--insecure** option means that this initial request will be without a certificate.

```
# curl --output hosted-engine.ca --insecure https://[Manager.example.com]/ca.crt
```

#### 3. Retrieving the GUID of the Host to be Removed

Use a **GET** request on the hosts collection to retrieve the Global Unique Identifier (GUID) for the host to be removed. The following example specifies this as a **GET** request, and includes the Manager CA certificate file. The **admin@internal** user is used for authentication, the password for which will be prompted once the command is executed.

```
# curl --request GET --cacert hosted-engine.ca --user admin@internal
https://[Manager.example.com]/api/hosts
```

This request will return information for all of the hosts in the environment. The host GUID is a hexadecimal string associated with the host name. For more information on the Red Hat Enterprise Virtualization REST API, see the *Red Hat Enterprise Virtualization Technical Guide*.

#### 4. Removing the Fenced Host

Use a **DELETE** request using the GUID of the fenced host to remove the host from the environment. In addition to the previously used options this example specifies headers to specify that the request is to be sent and returned using eXtensible Markup Language (XML), and the body in XML that sets the **force** action to be **true**.

■

```
curl --request DELETE --cacert hosted-engine.ca --user admin@internal --header "Content-Type: application/xml" --header "Accept: application/xml" --data "<action><force>true</force></action>" https://[Manager.example.com]/api/hosts/ecde42b0-de2f-48fe-aa23-1ebd5196b4a5
```

This **DELETE** request can be used to removed every fenced host in the self-hosted engine environment, as long as the appropriate GUID is specified.

Once the host has been removed, it can be re-installed to the self-hosted engine environment.

[Report a bug](#)

### 3.9.5. Installing Additional Hosts to a Restored Self-Hosted Engine Environment

Re-installing hosted-engine hosts that were present in a restored self-hosted engine environment at the time of the backup is slightly different to adding new hosts. Re-installing hosts will encounter the same VDSM timeout as the first host when it synchronized with the engine.

These hosts must have been removed from the environment before they can be re-installed.

As with the previous hosted-engine host, additional hosts require Red Hat Enterprise Linux 6.5, 6.6, or 7 with subscriptions to the appropriate Red Hat Enterprise Virtualization entitlements.

All steps in this procedure are to be conducted as the **root** user.

#### Procedure 3.9. Adding the host

1. Install the `ovirt-hosted-engine-setup` package.

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

2. Configure the host with the deployment command.

```
# hosted-engine --deploy
```

If running the **hosted-engine** deployment script over a network, it is recommended to use the **screen** window manager to avoid losing the session in case of network or terminal disruption. Install the `screen` package first if not installed.

```
# screen hosted-engine --deploy
```

3. **Configuring Storage**

Select the type of storage to use.

```
During customization use CTRL-D to abort.  
Please specify the storage you would like to use (iscsi, nfs3, nfs4)[nfs3]:
```

- For NFS storage types, specify the full address, using either the FQDN or IP address, and path name of the shared storage domain.

```
Please specify the full shared storage connection path to use (example: host:/path):  
storage.example.com:/hosted_engine/nfs
```

- For iSCSI, specify the iSCSI portal IP address, port, user name and password, and select a target name from the auto-detected list:

```
Please specify the iSCSI portal IP address:
Please specify the iSCSI portal port [3260]:
Please specify the iSCSI portal user:
Please specify the iSCSI portal password:
Please specify the target name (auto-detected values) [default]:
```

#### 4. Detecting the Self-Hosted Engine

The **hosted-engine** script detects that the shared storage is being used and asks if this is an additional host setup. You are then prompted for the host ID, which must be an integer not already assigned to a host in the environment.

```
The specified storage location already contains a data domain. Is this an additional host
setup (Yes, No)[Yes]?
[ INFO ] Installing on additional host
Please specify the Host ID [Must be integer, default: 2]:
```

#### 5. Configuring the System

The **hosted-engine** script uses the answer file generated by the original hosted-engine setup. To achieve this, the script requires the FQDN or IP address and the password of the **root** user of that host so as to access and secure-copy the answer file to the additional host.

```
[WARNING] A configuration file must be supplied to deploy Hosted Engine on an additional
host.
The answer file may be fetched from the first host using scp.
If you do not want to download it automatically you can abort the setup answering no to the
following question.
Do you want to scp the answer file from the first host? (Yes, No)[Yes]:
Please provide the FQDN or IP of the first host:
Enter 'root' user password for host [hosted_engine_1.example.com]:
[ INFO ] Answer file successfully downloaded
```

#### 6. Configuring the Hosted Engine

Specify the name for the additional host to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user. The name must not already be in use by a host in the environment.

```
Enter the name which will be used to identify this host inside the Administrator Portal
[hosted_engine_2]:
Enter 'admin@internal' user password that will be used for accessing the Administrator
Portal:
Confirm 'admin@internal' user password:
```

#### 7. Configuration Preview

Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface           : eth1
Engine FQDN                : HostedEngine-VM.example.com
Bridge name                : rhevm
SSH daemon port           : 22
Firewall manager           : iptables
```

```

Gateway address          : X.X.X.X
Host name for web application : hosted_engine_2
Host ID                  : 2
Image size GB            : 25
Storage connection       : storage.example.com:/hosted_engine/nfs
Console type             : vnc
Memory size MB           : 4096
MAC address              : 00:16:3e:05:95:50
Boot type                : disk
Number of CPUs           : 2
CPU Type                 : model_Penryn

```

Please confirm installation settings (Yes, No)[Yes]:

## 8. Confirming Engine Installation Complete

The additional host will contact the Manager and **hosted\_engine\_1**, after which the script will prompt for a selection. Continue by selection option 1.

```

[ INFO ] Stage: Closing up
        To continue make a selection from the options below:
        (1) Continue setup - engine installation is complete
        (2) Power off and restart the VM
        (3) Abort setup

(1, 2, 3)[1]:

```

9. [ INFO ] Engine replied: DB Up!Welcome to Health Status!  
 [ INFO ] Waiting for the host to become operational in the engine. This may take several minutes...

At this point, the host will become visible in the Administration Portal with **Installing** and **Initializing** states before entering a **Non Operational** state. The host will continue to wait for VDSM host to become operational until it eventually times out.

```

[ INFO ] Still waiting for VDSM host to become operational...
[ INFO ] Still waiting for VDSM host to become operational...
[ ERROR ] Timed out while waiting for host to start. Please check the logs.
[ ERROR ] Unable to add hosted_engine_1 to the manager
[ INFO ] Enabling and starting HA services
        Hosted Engine successfully set up
[ INFO ] Stage: Clean up
[ INFO ] Generating answer file '/etc/ovirt-hosted-engine/answers.conf'
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination

```

## 10. Activating the Host

- a. Log in to the Administration Portal.
- b. Click the **Hosts** tab and select the host to activate.
- c. Click the **Activate** button.

The host is now able to host the Manager virtual machine, and other virtual machines running in the self-hosted engine environment.

[Report a bug](#)

## 3.10. MIGRATING TO A SELF-HOSTED ENVIRONMENT

### Summary

Deploy a hosted-engine environment and migrate an existing instance of Red Hat Enterprise Virtualization. The **hosted-engine** deployment script is provided to assist with this task. The script asks you a series of questions, and configures your environment based on your answers. When the required values have been provided, the updated configuration is applied and the Red Hat Enterprise Virtualization Manager services are started.


The **hosted-engine** deployment script guides you through several distinct configuration stages. The script suggests possible configuration defaults in square brackets. Where these default values are acceptable, no additional input is required.

This procedure requires a new Red Hat Enterprise Linux 6.5 or 6.6 host with the `ovirt-hosted-engine-setup` package installed. This host is referred to as 'Host-HE1', with a fully qualified domain name (FQDN) of **Host-HE1.example.com** in this procedure.

Your original Red Hat Enterprise Virtualization Manager is referred to as 'BareMetal-Manager', with an FQDN of **Manager.example.com**, in this procedure. You are required to access and make changes on BareMetal-Manager during this procedure.

The hosted engine, the virtual machine created during configuration of Host-HE1 and used to manage the environment, is referred to as 'HostedEngine-VM' in this procedure. The **hosted-engine** deployment script prompts you to access this virtual machine multiple times to install an operating system and to configure the engine.

All steps in this procedure are to be conducted as the **root** user for the specified machine.



### IMPORTANT

The engine running on BareMetal-Manager must be the same version as will be installed on HostedEngine-VM. As the hosted engine feature is only available on Red Hat Enterprise Virtualization version 3.3.0 and later, any previous version of Red Hat Enterprise Virtualization running on BareMetal-Manager must be upgraded. Upgrade the engine version on BareMetal-Manager before creating the backup with the **engine-backup** command.

### Procedure 3.10. Migrating to a Self-Hosted Environment

#### 1. Initiating Hosted Engine Deployment

Begin configuration of the self-hosted environment by deploying the **hosted-engine** customization script on Host\_HE1. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

```
# hosted-engine --deploy
```

#### 2. Configuring Storage

Select the version of NFS and specify the full address, using either the FQDN or IP address, and path name of the shared storage domain. Choose the storage domain and storage data center names to be used in the environment.

During customization use CTRL-D to abort.

Please specify the storage you would like to use (nfs3, nfs4)[nfs3]:  
 Please specify the full shared storage connection path to use (example: host:/path):  
*storage.example.com:/hosted\_engine/nfs*  
 [ INFO ] Installing on first host  
 Please provide storage domain name. [hosted\_storage]:  
 Local storage datacenter name is an internal name and currently will not be shown in  
 engine's admin UI. Please enter local datacenter name [hosted\_datacenter]:

### 3. Configuring the Network

The script detects possible network interface controllers (NICs) to use as a management bridge for the environment. It then checks your firewall configuration and offers to modify it for console (SPICE or VNC) access HostedEngine-VM. Provide a pingable gateway IP address, to be used by the **ovirt-ha-agent** to help determine a host's suitability for running HostedEngine-VM.

Please indicate a nic to set rhevm bridge on: (eth1, eth0) [eth1]:  
 iptables was detected on your computer, do you wish setup to configure it? (Yes, No)[Yes]:  
 Please indicate a pingable gateway IP address [X.X.X.X]:


### 4. Configuring the Virtual Machine

The script creates a virtual machine to be configured as the Red Hat Enterprise Virtualization Manager, the hosted engine referred to in this procedure as HostedEngine-VM. Specify the boot device and, if applicable, the path name of the installation media, the CPU type, the number of virtual CPUs, and the disk size. Specify a MAC address for the HostedEngine-VM, or accept a randomly generated one. The MAC address can be used to update your DHCP server prior to installing the operating system on the virtual machine. Specify memory size and console connection type for the creation of HostedEngine-VM.

Please specify the device to boot the VM from (cdrom, disk, pxe) [cdrom]:  
 The following CPU types are supported by this host:  
 - model\_Penryn: Intel Penryn Family  
 - model\_Conroe: Intel Conroe Family  
 Please specify the CPU type to be used by the VM [model\_Penryn]:  
 Please specify the number of virtual CPUs for the VM [Defaults to minimum requirement: 2]:  
 Please specify the disk size of the VM in GB [Defaults to minimum requirement: 25]:  
 You may specify a MAC address for the VM or accept a randomly generated default  
 [00:16:3e:77:b2:a4]:  
 Please specify the memory size of the VM in MB [Defaults to minimum requirement: 4096]:  
 Please specify the console type you want to use to connect to the VM (vnc, spice) [vnc]:

### 5. Configuring the Hosted Engine

Specify the name for Host-HE1 to be identified in the Red Hat Enterprise Virtualization environment, and the password for the **admin@internal** user to access the Administrator Portal. Provide the FQDN for HostedEngine-VM; this procedure uses the FQDN *Manager.example.com*. Finally, 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.



#### IMPORTANT

The FQDN provided for the engine (*Manager.example.com*) must be the same FQDN provided when BareMetal-Manager was initially set up.

Enter the name which will be used to identify this host inside the Administrator Portal

```
[hosted_engine_1]: Host-HE1
Enter 'admin@internal' user password that will be used for accessing the Administrator
Portal:
Confirm 'admin@internal' user password:
Please provide the FQDN for the engine you want to use. This needs to match the FQDN that
you will use for the engine installation within the VM: Manager.example.com
Please provide the name of the SMTP server through which we will send notifications
[localhost]:
Please provide the TCP port number of the SMTP server [25]:
Please provide the email address from which notifications will be sent [root@localhost]:
Please provide a comma-separated list of email addresses which will get notifications
[root@localhost]:
```

## 6. Configuration Preview

Before proceeding, the **hosted-engine** script displays the configuration values you have entered, and prompts for confirmation to proceed with these values.

```
Bridge interface           : eth1
Engine FQDN                : Manager.example.com
Bridge name                : rhevm
SSH daemon port            : 22
Firewall manager           : iptables
Gateway address            : X.X.X.X
Host name for web application : Host-HE1
Host ID                    : 1
Image size GB              : 25
Storage connection         : storage.example.com:/hosted_engine/nfs
Console type               : vnc
Memory size MB             : 4096
MAC address                 : 00:16:3e:77:b2:a4
Boot type                   : pxe
Number of CPUs              : 2
CPU Type                    : model_Penryn

Please confirm installation settings (Yes, No)[No]:
```

## 7. Creating HostedEngine-VM

The script creates the virtual machine to be configured as HostedEngine-VM and provides connection details. You must install an operating system on HostedEngine-VM before the **hosted-engine** script can proceed on Host-HE1.

```
[ INFO ] Generating answer file '/etc/ovirt-hosted-engine/answers.conf'
[ INFO ] Stage: Transaction setup
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Configuring libvirt
[ INFO ] Generating VDSM certificates
[ INFO ] Configuring VDSM
[ INFO ] Starting vdsmd
[ INFO ] Waiting for VDSM hardware info
[ INFO ] Creating Storage Domain
[ INFO ] Creating Storage Pool
[ INFO ] Connecting Storage Pool
[ INFO ] Verifying sanlock lockspace initialization
```

```
[ INFO ] Initializing sanlock lockspace
[ INFO ] Initializing sanlock metadata
[ INFO ] Creating VM Image
[ INFO ] Disconnecting Storage Pool
[ INFO ] Start monitoring domain
[ INFO ] Configuring VM
[ INFO ] Updating hosted-engine configuration
[ INFO ] Stage: Transaction commit
[ INFO ] Stage: Closing up
[ INFO ] Creating VM
```

You can now connect to the VM with the following command:

```
/usr/bin/remote-viewer vnc://localhost:5900
```

Use temporary password "5379skAb" to connect to vnc console.

Please note that in order to use remote-viewer you need to be able to run graphical applications.

This means that if you are using ssh you have to supply the -Y flag (enables trusted X11 forwarding).

Otherwise you can run the command from a terminal in your preferred desktop environment.

If you cannot run graphical applications you can connect to the graphic console from another host or connect to the console using the following command:

```
virsh -c qemu+tls://Test/system console HostedEngine
```

If you need to reboot the VM you will need to start it manually using the command:

```
hosted-engine --vm-start
```

You can then set a temporary password using the command:

```
hosted-engine --add-console-password
```

The VM has been started. Install the OS and shut down or reboot it. To continue please make a selection:

- (1) Continue setup - VM installation is complete
- (2) Reboot the VM and restart installation
- (3) Abort setup

```
(1, 2, 3)[1]:
```

Using the naming convention of this procedure, connect to the virtual machine using VNC with the following command:

```
/usr/bin/remote-viewer vnc://Host-HE1.example.com:5900
```

## 8. Installing the Virtual Machine Operating System

Connect to HostedEngine-VM, the virtual machine created by the hosted-engine script, and install a Red Hat Enterprise Linux 6.5 or 6.6 operating system.

## 9. Synchronizing the Host and the Virtual Machine

Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

- (1) Continue setup - VM installation is complete

```
Waiting for VM to shut down...
```

```
[ INFO ] Creating VM
```

You can now connect to the VM with the following command:

```
/usr/bin/remote-viewer vnc://localhost:5900
```

Use temporary password "5379skAb" to connect to vnc console.

Please note that in order to use remote-viewer you need to be able to run graphical applications.



This means that if you are using ssh you have to supply the `-Y` flag (enables trusted X11 forwarding).

Otherwise you can run the command from a terminal in your preferred desktop environment. If you cannot run graphical applications you can connect to the graphic console from another host or connect to the console using the following command:

```
virsh -c qemu+tls://Test/system console HostedEngine
```

If you need to reboot the VM you will need to start it manually using the command:

```
hosted-engine --vm-start
```

You can then set a temporary password using the command:

```
hosted-engine --add-console-password
```

Please install and setup the engine in the VM.

You may also be interested in subscribing to "agent" RHN/Satellite channel and installing `rhev-guest-agent-common` package in the VM.

To continue make a selection from the options below:

- (1) Continue setup - engine installation is complete
- (2) Power off and restart the VM
- (3) Abort setup

## 10. Installing the Manager

Connect to HostedEngine-VM, subscribe to the appropriate Red Hat Enterprise Virtualization Manager channels, ensure that the most up-to-date versions of all installed packages are in use, and install the `rhev` packages.

```
# yum upgrade
```

```
# yum install rhevm
```

## 11. Disabling BareMetal-Manager

Connect to BareMetal-Manager, the Manager of your established Red Hat Enterprise Virtualization environment, and stop the engine and prevent it from running.

```
# service ovirt-engine stop
# service ovirt-engine disable
# chkconfig ovirt-engine off
```



### NOTE

Though stopping BareMetal-Manager from running is not obligatory, it is recommended as it ensures no changes will be made to the environment after the backup has been created. Additionally, it prevents BareMetal-Manager and HostedEngine-VM from simultaneously managing existing resources.

## 12. Updating DNS

Update your DNS so that the FQDN of the Red Hat Enterprise Virtualization environment correlates to the IP address of HostedEngine-VM and the FQDN previously provided when configuring the **hosted-engine** deployment script on Host-HE1. In this procedure, FQDN was set as *Manager.example.com* because in a migrated hosted-engine setup, the FQDN provided for the engine must be identical to that given in the engine setup of the original engine.

## 13. Creating a Backup of BareMetal-Manager

Connect to BareMetal-Manager and run the **engine-backup** command with the `--mode=backup`, `--file=[FILE]`, and `--log=[LogFILE]` parameters to specify the backup mode, the name of the backup file created and used for the backup, and the name of the log file to be

created to store the backup log.

```
# engine-backup --mode=backup --file=[FILE] --log=[LogFILE]
```

#### 14. Copying the Backup File to HostedEngine-VM

On BareMetal-Manager, secure copy the backup file to HostedEngine-VM. In the following example, *[Manager.example.com]* is the FQDN for HostedEngine-VM, and */backup/* is any designated folder or path. If the designated folder or path does not exist, you must connect to HostedEngine-VM and create it before secure copying the backup from BareMetal-Manager.

```
# scp -p backup1 [Manager.example.com:/backup/]
```

#### 15. Restoring the Backup File on HostedEngine-VM

The **engine-backup --mode=restore** command does not create a database; you are required to create one on HostedEngine-VM before restoring the backup you created on BareMetal-Manager. Connect to HostedEngine-VM and create the database, as detailed in [Section 2.4.4, “Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager”](#).



#### NOTE

The procedure in [Section 2.4.4, “Preparing a PostgreSQL Database for Use with Red Hat Enterprise Virtualization Manager”](#) creates a database that is not empty, which will result in the following error when you attempt to restore the backup:

```
FATAL: Database is not empty
```

Create an empty database using the following command in psql:

```
postgres=# create database [database name] owner [user name]
```

After the empty database has been created, restore the BareMetal-Manager backup using the **engine-backup** command with the **--mode=restore --file=[FILE] --log=[Restore.log]** parameters to specify the restore mode, the name of the file to be used to restore the database, and the name of the logfile to store the restore log. This restores the files and the database but does not start the service.

To specify a different database configuration, use the **--change-db-credentials** parameter to activate alternate credentials. Use the **engine-backup --help** command on the Manager for a list of credential parameters.

```
# engine-backup --mode=restore --file=[FILE] --log=[Restore.log] --change-db-credentials --db-host=[X.X.X.X] --db-user=[engine] --db-password=[password] --db-name=[engine]
```

#### 16. Configuring HostedEngine-VM

Configure the engine on HostedEngine-VM. This identifies the existing files and database.

```
# engine-setup
```

```
[ INFO ] Stage: Initializing
```

```
[ INFO ] Stage: Environment setup
```

```
Configuration files: [/etc/ovirt-engine-setup.conf.d/10-packaging.conf, '/etc/ovirt-engine-setup.conf.d/20-setup-ovirt-post.conf']
```

Log file: /var/log/ovirt-engine/setup/ovirt-engine-setup-20140304075238.log

Version: otopi-1.1.2 (otopi-1.1.2-1.el6ev)

[ INFO ] Stage: Environment packages setup

[ INFO ] Yum Downloading: rhel-65-zstream/primary\_db 2.8 M(70%)

[ INFO ] Stage: Programs detection

[ INFO ] Stage: Environment setup

[ INFO ] Stage: Environment customization

==== PACKAGES ====

[ INFO ] Checking for product updates...

[ INFO ] No product updates found

==== NETWORK CONFIGURATION ====

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]:

[ INFO ] iptables will be configured as firewall manager.

==== DATABASE CONFIGURATION ====

==== OVIRT ENGINE CONFIGURATION ====

Skipping storing options as database already prepared

==== PKI CONFIGURATION ====

PKI is already configured

==== APACHE CONFIGURATION ====

==== SYSTEM CONFIGURATION ====

==== END OF CONFIGURATION ====

[ INFO ] Stage: Setup validation

[WARNING] Less than 16384MB of memory is available

[ INFO ] Cleaning stale zombie tasks

==== CONFIGURATION PREVIEW ====

Database name : *engine*  
Database secured connection : False  
Database host : *X.X.X.X*  
Database user name : *engine*  
Database host name validation : False  
Database port : 5432  
NFS setup : True  
Firewall manager : iptables  
Update Firewall : True  
Configure WebSocket Proxy : True  
Host FQDN : *Manager.example.com*

```
NFS mount point           : /var/lib/exports/iso
Set application as default page  : True
Configure Apache SSL       : True
```

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

Confirm the settings. Upon completion, the setup provides an SSH fingerprint and an internal Certificate Authority hash.

## 17. Synchronizing the Host and the Manager

Return to Host-HE1 and continue the **hosted-engine** deployment script by selecting option 1:

```
(1) Continue setup - engine installation is complete
```

```
[ INFO ] Engine replied: DB Up!Welcome to Health Status!
[ INFO ] Waiting for the host to become operational in the engine. This may take several
minutes...
[ INFO ] Still waiting for VDSM host to become operational...
[ INFO ] The VDSM Host is now operational
        Please shutdown the VM allowing the system to launch it as a monitored service.
        The system will wait until the VM is down.
```

## 18. Shutting Down HostedEngine-VM

Shutdown HostedEngine-VM.

```
# shutdown now
```

## 19. Setup Confirmation

Return to Host-HE1 to confirm it has detected that HostedEngine-VM is down.

```
[ INFO ] Enabling and starting HA services
        Hosted Engine successfully set up
[ INFO ] Stage: Clean up
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination
```

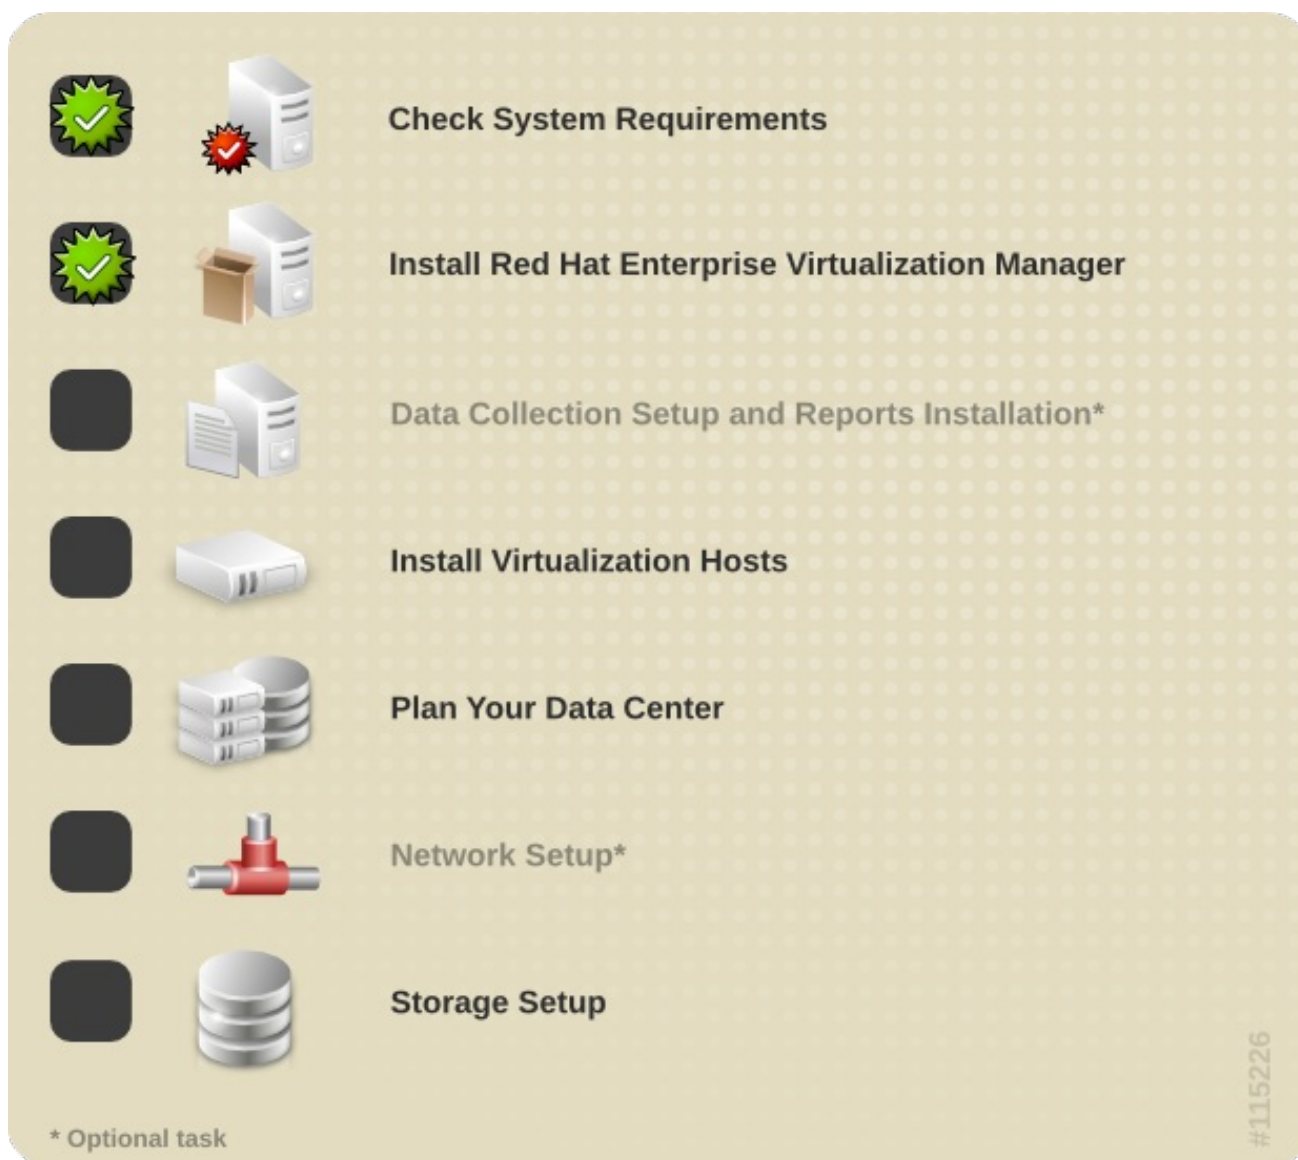
## Result

Your Red Hat Enterprise Virtualization engine has been migrated to a hosted-engine setup. The Manager is now operating on a virtual machine on Host-HE1, called HostedEngine-VM in the environment. As HostedEngine-VM is highly available, it is migrated to other hosts in the environment when applicable.

[Report a bug](#)

## CHAPTER 4. HISTORY AND REPORTS

### 4.1. WORKFLOW PROGRESS - DATA COLLECTION SETUP AND REPORTS INSTALLATION



[Report a bug](#)

### 4.2. DATA COLLECTION SETUP AND REPORTS INSTALLATION OVERVIEW

The Red Hat Enterprise Virtualization Manager optionally includes a comprehensive management history database, which can be utilized by any application to extract a range of information at the data center, cluster, and host levels. As the database structure changes over time a number of database views are also included to provide a consistent structure to consuming applications. A view is a virtual table composed of the result set of a database query. The definition of a view is stored in the database as a **SELECT** statement. The result set of the **SELECT** statement populates the virtual table returned by the view. If the optional comprehensive management history database has been enabled, the history tables and their associated views are stored in the **ovirt\_engine\_history** database.

In addition to the history database Red Hat Enterprise Virtualization Manager Reports functionality is

also available as an optional component. Red Hat Enterprise Virtualization Manager Reports provides a customized implementation of JasperServer, and JasperReports. JasperServer is a component of JasperReports, an open source reporting tool capable of being embedded in Java-based applications. It produces reports which can be rendered to screen, printed, or exported to a variety of formats including PDF, Excel, CSV, Word, RTF, Flash, ODT and ODS. Reports built in Red Hat Enterprise Virtualization Manager Reports are accessed via a web interface. In addition to a range of pre-configured reports and dashboards for monitoring the system, you are also able to create your own ad hoc reports.

Before proceeding with Red Hat Virtualization Manager Reports installation you must first have installed the Red Hat Enterprise Virtualization Manager.

The Red Hat Enterprise Virtualization Manager Reports functionality depends on the presence of the history database, which is installed separately. Both the history database and the Red Hat Enterprise Virtualization Manager Reports are optional components. They are not installed by default when you install the Red Hat Enterprise Virtualization Manager.



#### NOTE


Detailed user, administration, and installation guides for JasperReports can be found in </usr/share/jasperreports-server-pro/docs/>

[Report a bug](#)

## 4.3. INSTALLING AND CONFIGURING THE HISTORY DATABASE AND RED HAT ENTERPRISE VIRTUALIZATION MANAGER REPORTS

### Summary

Use of the history database and reports is optional. To use the reporting capabilities of Red Hat Enterprise Virtualization Manager, you must install and configure `rhev-dwh` and `rhev-reports`.



#### IMPORTANT

If you are using the Self-Hosted Engine, you must move it to maintenance mode:

```
# hosted-engine --set-maintenance --mode=global
```

### Procedure 4.1. Installing and Configuring the History Database and Red Hat Enterprise Virtualization Manager Reports

1. Install the `rhev-dwh` package. This package must be installed on the system on which the Red Hat Enterprise Virtualization Manager is installed.

```
# yum install rhev-dwh
```

2. Install the `rhev-reports` package. This package must be installed on the system on which the Red Hat Enterprise Virtualization Manager is installed.

```
# yum install rhev-reports
```

3. Run the **engine-setup** command on the system hosting the Red Hat Enterprise Virtualization Manager and follow the prompts to configure Data Warehouse and Reports:

```
--== PRODUCT OPTIONS ==--
```

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

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

- The command will prompt you to answer the following questions about the DWH database:

```
--== DATABASE CONFIGURATION ==--
```

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

```
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]:
```

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

```
Setup can configure the local postgresql server automatically for the Reports to run. This may conflict with existing applications.
```

```
Would you like Setup to automatically configure postgresql and create Reports database, or prefer to perform that manually? (Automatic, Manual) [Automatic]:
```

Press **Enter** to choose the highlighted defaults, or type your alternative preference then press **Enter**.

- The command will then prompt you to set the password for the Red Hat Enterprise Virtualization Manager Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials which are separate to those used for Red Hat Enterprise Virtualization Manager.

```
Reports power users password:
```

You will be prompted to enter the password a second time to confirm it.

- For the Red Hat Enterprise Virtualization Manager Reports installation to take effect, the **ovirt-engine** service must be restarted. The **engine-setup** command prompts you:


```
During execution engine service will be stopped (OK, Cancel) [OK]:
```

Type **OK** and then press **Enter** to proceed. The **ovirt-engine** service will restart automatically later in the command.

## Result

The **ovirt\_engine\_history** database has been created. Red Hat Enterprise Virtualization Manager is configured to log information to this database for reporting purposes. Red Hat Enterprise Virtualization Manager Reports has been installed successfully. Access Red Hat Enterprise Virtualization Manager Reports at [http://\[demo.redhat.com\]/ovirt-engine-reports](http://[demo.redhat.com]/ovirt-engine-reports), replacing **[demo.redhat.com]** with the fully-qualified domain name of the Red Hat Enterprise Virtualization Manager. If during Red Hat Enterprise Virtualization Manager installation you selected a non-default HTTP port then append **:[port]** to the URL, replacing **[port]** with the port that you chose.

Use the user name **admin** and the password you set during reports installation to log in for the first time. Note that the first time you log into Red Hat Enterprise Virtualization Manager Reports, a number of web pages are generated, and as a result your initial attempt to login may take some time to complete.



## NOTE

Previously, the **admin** user name was **rhev-admin**. If you are performing a clean installation, the user name is now **admin**. In you are performing an upgrade, the user name will remain **rhev-admin**.

[Report a bug](#)



# CHAPTER 5. UPDATING THE RED HAT ENTERPRISE VIRTUALIZATION ENVIRONMENT

## 5.1. UPGRADES BETWEEN MINOR RELEASES

### 5.1.1. Checking for Red Hat Enterprise Virtualization Manager Updates



#### IMPORTANT

Always update to the latest minor version of your current Red Hat Enterprise Virtualization Manager version before you upgrade to the next major version.

#### Summary

Check for updates to the Red Hat Enterprise Virtualization Manager.

#### Procedure 5.1. Checking for Red Hat Enterprise Virtualization Manager Updates

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

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

2.
  - If there are no updates available, the command will output the text **No upgrade:**

```
# engine-upgrade-check
VERB: queue package rhvm-setup for update
VERB: package rhvm-setup queued
VERB: Building transaction
VERB: Empty transaction
VERB: Transaction Summary:
No upgrade
```

- If updates are available, the command will list the packages to be updated:

```
# engine-upgrade-check
VERB: queue package rhvm-setup for update
VERB: package rhvm-setup queued
VERB: Building transaction
VERB: Transaction built
VERB: Transaction Summary:
VERB:  updated - rhvm-lib-3.3.2-0.50.el6ev.noarch
VERB:  update  - rhvm-lib-3.4.0-0.13.el6ev.noarch
VERB:  updated - rhvm-setup-3.3.2-0.50.el6ev.noarch
VERB:  update  - rhvm-setup-3.4.0-0.13.el6ev.noarch
VERB:  install - rhvm-setup-base-3.4.0-0.13.el6ev.noarch
VERB:  install - rhvm-setup-plugin-ovirt-engine-3.4.0-0.13.el6ev.noarch
VERB:  updated - rhvm-setup-plugins-3.3.1-1.el6ev.noarch
VERB:  update  - rhvm-setup-plugins-3.4.0-0.5.el6ev.noarch
```

Upgrade available

Upgrade available

## Result

You have checked for updates to the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 5.1.2. Updating the Red Hat Enterprise Virtualization Manager

### Summary

Updates to the Red Hat Enterprise Virtualization Manager are released via Red Hat Network. Before installing an update from Red Hat Network, ensure you read the advisory text associated with it and the latest version of the *Red Hat Enterprise Virtualization Release Notes* and *Red Hat Enterprise Virtualization Technical Notes*. A number of actions must be performed to complete an upgrade, including:

- Stopping the **ovirt-engine** service.
- Downloading and installing the updated packages.
- Backing up and updating the database.
- Performing post-installation configuration.
- Starting the **ovirt-engine** service.


### Procedure 5.2. Updating Red Hat Enterprise Virtualization Manager

1. Run the following command to update the `rhevms-setup` package:

```
# yum update rhevm-setup
```

2. Run the following command to update the Red Hat Enterprise Virtualization Manager:

```
# engine-setup
```



#### IMPORTANT

Active hosts are not updated by this process and must be updated separately. As a result, the virtual machines running on those hosts are not affected.



#### IMPORTANT

The update process may take some time; allow time for the update process to complete and do not stop the process once initiated. Once the update is complete, you will also be instructed to separately update the data warehouse and reports functionality. These additional steps are only required if you installed these features.

## Result

You have successfully updated the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

### 5.1.3. Updating Red Hat Enterprise Virtualization Hypervisors

#### Summary


Updating Red Hat Enterprise Virtualization Hypervisors involves reinstalling the Hypervisor with a newer version of the Hypervisor ISO image. This includes stopping and restarting the Hypervisor. Virtual machines are automatically migrated to a different host, as a result it is recommended that Hypervisor updates are performed at a time when the host's usage is relatively low.

It is recommended that administrators update Red Hat Enterprise Virtualization Hypervisors regularly. Important bug fixes and security updates are included in updates. Hypervisors which are not up to date may be a security risk.



#### WARNING

Upgrading Hypervisor hosts involves shutting down, deactivating guests, and restarting the physical server. If any virtual machines are running on the Hypervisor, all data and configuration details may be destroyed if they are not shut down. Upgrading Hypervisors must be carefully planned and executed with care and consideration.



#### IMPORTANT

Ensure that the cluster contains more than one host before performing an upgrade. Do not attempt to reinstall or upgrade all the hosts at the same time, as one host must remain available to perform Storage Pool Manager (SPM) tasks.

#### Procedure 5.3. Updating Red Hat Enterprise Virtualization Hypervisors

1. Log in to the system hosting Red Hat Enterprise Virtualization Manager as the **root** user.
2. Enable the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86\_64)** repository:

- With RHN Classic:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev
```

- With Subscription Manager, attach a **Red Hat Enterprise Virtualization** entitlement and run the following command:

```
# subscription-manager repos --enable=rhel-6-server-rhev-rpms
```

3. Run the **yum** command with the **update rhev-hypervisor6** parameters to ensure that you have the most recent version of the rhev-hypervisor6 package installed.

```
# yum update rhel-hypervisor6
```

4. Use your web browser to log in to the Administration Portal as a Red Hat Enterprise Virtualization administrative user.
5. Click the **Hosts** tab, and then select the host that you intend to upgrade. If the host is not displayed, or the list of hosts is too long to filter visually, perform a search to locate the host.
6. With the host selected, click the **General** tab in the details pane.
  - If the host requires updating, an alert message indicates that a new version of the Red Hat Enterprise Virtualization Hypervisor is available.
  - If the host does not require updating, no alert message is displayed and no further action is required.
7. Ensure the host remains selected and click the **Maintenance** button, if the host is not already in maintenance mode. This will cause any virtual machines running on the host to be migrated to other hosts. If the host is the SPM, this function will be moved to another host. The status of the host changes as it enters maintenance mode. When the host status is **Maintenance**, the message in the general tab changes, providing you with a link which when clicked will reinstall or upgrade the host.
8. Ensure that the host remains selected, and that you are on the **General** tab of the details pane. Click the **Upgrade** link to open the **Install Host** window.
9. Select **rhel-hypervisor.iso**, which is symbolically linked to the most recent hypervisor image.
10. Click **OK** to update and reinstall the host. The dialog closes, the details of the host are updated in the **Hosts** tab, and the status changes.

The host status will transition through these stages:

- **Installing**
- **Reboot**
- **Non Responsive**
- **Up.**

These are all expected, and each stage will take some time.

11. Once successfully updated, the host displays a status of **Up**. Any virtual machines that were migrated off the host, are at this point able to be migrated back to it.



## IMPORTANT

After a Red Hat Enterprise Virtualization Hypervisor is successfully registered to the Red Hat Enterprise Virtualization Manager and then upgraded, it may erroneously appear in the Administration Portal with the status of **Install Failed**. Click on the **Activate** button, and the hypervisor will change to an **Up** status and be ready for use.

## Result

You have successfully updated a Red Hat Enterprise Virtualization Hypervisor. Repeat these steps for each Hypervisor in the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

## 5.1.4. Updating Red Hat Enterprise Linux Virtualization Hosts

### Summary

Red Hat Enterprise Linux hosts are using the **yum** in the same way as regular Red Hat Enterprise Linux systems. It is highly recommended that you use **yum** to update your systems regularly, to ensure timely application of security and bug fixes.

### Procedure 5.4. Updating Red Hat Enterprise Linux Hosts

1. From the Administration Portal, click the **Hosts** tab and select the host to be updated. Click **Maintenance** to place it into maintenance mode.
2. On the Red Hat Enterprise Linux host, run the following command:

```
# yum update
```

3. Restart the host to ensure all updates are correctly applied.

### Result

You have successfully updated the Red Hat Enterprise Linux host. Repeat this process for each Red Hat Enterprise Linux host in the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

## 5.1.5. Updating the Red Hat Enterprise Virtualization Guest Tools

### Summary

The guest tools comprise software that allows Red Hat Enterprise Virtualization Manager to communicate with the virtual machines it manages, providing information such as the IP addresses, memory usage, and applications installed on those virtual machines. The guest tools are distributed as an ISO file that can be attached to guests. This ISO file is packaged as an RPM file that can be installed and upgraded from the machine on which the Red Hat Enterprise Virtualization Manager is installed.

### Procedure 5.5. Updating the Red Hat Enterprise Virtualization Guest Tools

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

```
# yum update -y rhel-guest-tools-iso*
```

2. Run the following command to upload the ISO file to your ISO domain, replacing *[ISODomain]* with the name of your ISO domain:

```
engine-iso-uploader --iso-domain=[ISODomain] upload /usr/share/rhel-guest-tools-iso/rhel-guest-tools-setup.iso
```

**NOTE**

The **rhev-tools-setup.iso** file is a symbolic link to the most recently updated ISO file. The link is automatically changed to point to the newest ISO file every time you update the `rhev-guest-tools-iso` package.

3. Use the Administration Portal, User Portal, or REST API to attach the **rhev-tools-setup.iso** file to each of your virtual machines and upgrade the tools installed on each guest using the installation program on the ISO.

**Result**

You have updated the **rhev-tools-setup.iso** file, uploaded the updated ISO file to your ISO domain, and attached it to your virtual machines.

[Report a bug](#)

## 5.2. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION 3.4

### 5.2.1. Red Hat Enterprise Virtualization Manager 3.4 Upgrade Overview

**IMPORTANT**

Always update to the latest minor version of your current Red Hat Enterprise Virtualization Manager version before you upgrade to the next major version.

The process for upgrading Red Hat Enterprise Virtualization Manager comprises three main steps:

- Configuring channels and entitlements.
- Updating the required packages.
- Performing the upgrade.

The command used to perform the upgrade itself is **engine-setup**, which provides an interactive interface. While the upgrade is in process, virtualization hosts and the virtual machines running on those virtualization hosts continue to operate independently. When the upgrade is complete, you can then upgrade your hosts to the latest versions of Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor.

[Report a bug](#)

### 5.2.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.4

Some of the features provided by Red Hat Enterprise Virtualization 3.4 are only available if your data centers, clusters, and storage have a compatibility version of 3.4.

**Table 5.1. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.4**

Feature	Description
---------	-------------

Feature	Description
Abort migration on error	This feature adds support for handling errors encountered during the migration of virtual machines.
Forced Gluster volume creation	This feature adds support for allowing the creation of Gluster bricks on root partitions. With this feature, you can choose to override warnings against creating bricks on root partitions.
Management of asynchronous Gluster volume tasks	This feature provides support for managing asynchronous tasks on Gluster volumes, such as rebalancing volumes or removing bricks. To use this feature, you must use GlusterFS version 3.5 or above.
Import Glance images as templates	This feature provides support for importing images from an OpenStack image service as templates.
File statistic retrieval for non-NFS ISO domains	This feature adds support for retrieving statistics on files stored in ISO domains that use a storage format other than NFS, such as a local ISO domain.
Default route support	This feature adds support for ensuring that the default route of the management network is registered in the main routing table and that registration of the default route for all other networks is disallowed. This ensures the management network gateway is set as the default gateway for hosts.
Virtual machine reboot	This feature adds support for rebooting virtual machines from the User Portal or Administration Portal via a new button. To use this action on a virtual machine, you must install the guest tools on that virtual machine.

[Report a bug](#)

### 5.2.3. Red Hat Enterprise Virtualization 3.4 Upgrade Considerations

The following is a list of key considerations that must be made when planning your upgrade.



## IMPORTANT

### Upgrading to version 3.4 can only be performed from version 3.3

To upgrade a previous version of Red Hat Enterprise Virtualization earlier than Red Hat Enterprise Virtualization 3.3 to Red Hat Enterprise Virtualization 3.4, you must sequentially upgrade to any newer versions of Red Hat Enterprise Virtualization before upgrading to the latest version. For example, if you are using Red Hat Enterprise Virtualization 3.2, you must upgrade to Red Hat Enterprise Virtualization 3.3 before you can upgrade to Red Hat Enterprise Virtualization 3.4.

### Red Hat Enterprise Virtualization Manager cannot be installed on the same machine as IPA

An error message displays if the ipa-server package is installed. Red Hat Enterprise Virtualization Manager 3.4 does not support installation on the same machine as Identity Management (IdM). To resolve this issue, you must migrate the IdM configuration to another system before re-attempting the upgrade.

### Upgrading to JBoss Enterprise Application Platform 6.2 is recommended

Although Red Hat Enterprise Virtualization Manager 3.4 supports Enterprise Application Platform 6.1.0, upgrading to the latest supported version of JBoss is recommended.

[Report a bug](#)

## 5.2.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.4

### Summary

The following procedure outlines the process for upgrading Red Hat Enterprise Virtualization Manager 3.3 to Red Hat Enterprise Virtualization Manager 3.4. This procedure assumes that the system on which the Manager is installed is subscribed to the channels and entitlements for receiving Red Hat Enterprise Virtualization 3.3 packages at the start of the procedure.



## IMPORTANT

If the upgrade fails, the **engine-setup** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. For this reason, the channels required by Red Hat Enterprise Virtualization 3.3 must not be removed until after the upgrade is complete as outlined below. If the upgrade fails, detailed instructions display that explain how to restore your installation.

### Procedure 5.6. Upgrading to Red Hat Enterprise Virtualization Manager 3.4

1. Subscribe the system on which the Red Hat Enterprise Virtualization Manager is installed to the required channels and entitlements for receiving Red Hat Enterprise Virtualization Manager 3.4 packages.
  - With RHN Classic:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.4
```

- With Subscription Manager:

```
# yum-config-manager --enable rhel-6-server-rhev-3.4-rpms
```



2. Run the following command to ensure you have the most recent version of `engine-setup` by updating the `rhev-setup` package.

```
# yum update rhvm-setup
```

3. If you have installed Reports and the Data Warehouse, run the following command to ensure you have the most recent version of the `rhev-reports-setup` and `rhev-dwh-setup` packages:

```
# yum install rhvm-reports-setup rhvm-dwh-setup
```

4. Run the following command and follow the prompts to upgrade the Red Hat Enterprise Virtualization Manager:

```
# engine-setup
```

5. Remove or disable the Red Hat Enterprise Virtualization Manager 3.3 channel to ensure the system does not use any Red Hat Enterprise Virtualization Manager 3.3 packages.

- With RHN Classic:

```
# rhn-channel --remove --channel=rhel-x86_64-server-6-rhev-3.3
```

- With Subscription Manager:

```
# yum-config-manager --disable rhel-6-server-rhev-3.3-rpms
```

6. Run the following command to ensure all packages are up to date:

```
# yum update
```

## Result

You have upgraded the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 5.3. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION 3.3

### 5.3.1. Red Hat Enterprise Virtualization Manager 3.3 Upgrade Overview

Upgrading Red Hat Enterprise Virtualization Manager is a straightforward process that comprises three main steps:

- Configuring channels and entitlements.
- Updating the required packages.
- Performing the upgrade.

The command used to perform the upgrade itself is **engine-setup**, which provides an interactive interface. While the upgrade is in process, virtualization hosts and the virtual machines running on those virtualization hosts continue to operate independently. When the upgrade is complete, you can then

upgrade your hosts to the latest versions of Red Hat Enterprise Linux or Red Hat Enterprise Virtualization Hypervisor.

[Report a bug](#)

### 5.3.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3

Some of the new features in Red Hat Enterprise Virtualization are only available if your data centers, clusters, and storage have a compatibility version of **3.3**.

**Table 5.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3**

Feature	Description
Libvirt-to-libvirt virtual machine migration	Perform virtual machine migration using libvirt-to-libvirt communication. This is safer, more secure, and has less host configuration requirements than native KVM migration, but has a higher overhead on the host CPU.
Isolated network to carry virtual machine migration traffic	Separates virtual machine migration traffic from other traffic types, like management and display traffic. Reduces chances of migrations causing a network flood that disrupts other important traffic types.
Define a gateway per logical network	Each logical network can have a gateway defined as separate from the management network gateway. This allows more customizable network topologies.
Snapshots including RAM	Snapshots now include the state of a virtual machine's memory as well as disk.
Optimized iSCSI device driver for virtual machines	Virtual machines can now consume iSCSI storage as virtual hard disks using an optimized device driver.
Host support for MOM management of memory overcommitment	MOM is a policy-driven tool that can be used to manage overcommitment on hosts. Currently MOM supports control of memory ballooning and KSM.
GlusterFS data domains.	Native support for the GlusterFS protocol was added as a way to create storage domains, allowing Gluster data centers to be created.
Custom device property support	In addition to defining custom properties of virtual machines, you can also define custom properties of virtual machine devices.
Multiple monitors using a single virtual PCI device	Drive multiple monitors using a single virtual PCI device, rather than one PCI device per monitor.
Updatable storage server connections	It is now possible to edit the storage server connection details of a storage domain.

Feature	Description
Check virtual hard disk alignment	Check if a virtual disk, the filesystem installed on it, and its underlying storage are aligned. If it is not aligned, there may be a performance penalty.
Extendable virtual machine disk images	You can now grow your virtual machine disk image when it fills up.
OpenStack Image Service integration	Red Hat Enterprise Virtualization supports the OpenStack Image Service. You can import images from and export images to an Image Service repository.
Gluster hook support	You can manage Gluster hooks, which extend volume life cycle events, from Red Hat Enterprise Virtualization Manager.
Gluster host UUID support	This feature allows a Gluster host to be identified by the Gluster server UUID generated by Gluster in addition to identifying a Gluster host by IP address.
Network quality of service (QoS) support	Limit the inbound and outbound network traffic at the virtual NIC level.
Cloud-Init support	Cloud-Init allows you to automate early configuration tasks in your virtual machines, including setting hostnames, authorized keys, and more.

[Report a bug](#)

### 5.3.3. Red Hat Enterprise Virtualization 3.3 Upgrade Considerations

The following is a list of key considerations that must be made when planning your upgrade.



## IMPORTANT

### Upgrading to version 3.3 can only be performed from version 3.2

Users of Red Hat Enterprise Virtualization 3.1 must migrate to Red Hat Enterprise Virtualization 3.2 before attempting to upgrade to Red Hat Enterprise Virtualization 3.3.

### Red Hat Enterprise Virtualization Manager cannot be installed on the same machine as IPA

An error message displays if the ipa-server package is installed. Red Hat Enterprise Virtualization Manager 3.3 does not support installation on the same machine as Identity Management (IdM). To resolve this issue, you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information, see <https://access.redhat.com/knowledge/articles/233143>.

Error: IPA was found to be installed on this machine. Red Hat Enterprise Virtualization Manager 3.3 does not support installing IPA on the same machine. Please remove ipa packages before you continue.

### Upgrading to JBoss Enterprise Application Platform 6.1.0 is recommended

Although Red Hat Enterprise Virtualization Manager 3.3 supports Enterprise Application Platform 6.0.1, upgrading to the latest supported version of JBoss is recommended. For more information on upgrading to JBoss Enterprise Application Platform 6.1.0, see [Upgrade the JBoss EAP 6 RPM Installation](#).

### The rhvm-upgrade command has been replaced by engine-setup

From Version 3.3, installation of Red Hat Enterprise Virtualization Manager supports **otopi**, a standalone, plug-in-based installation framework for setting up system components. Under this framework, the **rhvm-upgrade** command used during the installation process has been updated to **engine-setup** and is now obsolete.

[Report a bug](#)

## 5.3.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.3

### Summary

The following procedure outlines the process for upgrading Red Hat Enterprise Virtualization Manager 3.2 to Red Hat Enterprise Virtualization Manager 3.3. This procedure assumes that the system on which the Manager is hosted is subscribed to the channels and entitlements for receiving Red Hat Enterprise Virtualization 3.2 packages.

If the upgrade fails, the **engine-setup** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. For this reason, the channels required by Red Hat Enterprise Virtualization 3.2 must not be removed until after the upgrade is complete as outlined below. If the upgrade fails, detailed instructions display that explain how to restore your installation.

### Procedure 5.7. Upgrading to Red Hat Enterprise Virtualization Manager 3.3

1. Subscribe the system to the required channels and entitlements for receiving Red Hat Enterprise Virtualization Manager 3.3 packages.

### Subscription Manager

Red Hat Enterprise Virtualization 3.3 packages are provided by the **rhel-6-server-rhev-3.3-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration.

```
# yum-config-manager --enable rhel-6-server-rhev-3.3-rpms
```

### Red Hat Network Classic

The Red Hat Enterprise Virtualization 3.3 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.3 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.3** in Red Hat Network Classic. Use the **rhn-channel** command or the Red Hat Network web interface to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.3 x86\_64)** channel:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.3
```

2. Update the **rhev-setup** package to ensure you have the most recent version of **engine-setup**.

```
# yum update rhev-setup
```

3. Run the **engine-setup** command and follow the prompts to upgrade Red Hat Enterprise Virtualization Manager.

```
# engine-setup
[ INFO ] Stage: Initializing
```

```
Welcome to the RHEV 3.3.0 upgrade.
```

```
Please read the following knowledge article for known issues and
updated instructions before proceeding with the upgrade.
```

```
RHEV 3.3.0 Upgrade Guide: Tips, Considerations and Roll-back Issues
```

```
https://access.redhat.com/site/articles/408623
```

```
Would you like to continue with the upgrade? (Yes, No) [Yes]:
```

4. Remove Red Hat Enterprise Virtualization Manager 3.2 channels and entitlements to ensure the system does not use any Red Hat Enterprise Virtualization Manager 3.2 packages.

### Subscription Manager

Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.2 repository in your **yum** configuration.

```
# yum-config-manager --disable rhel-6-server-rhev-3.2-rpms
```

### Red Hat Network Classic

Use the **rhn-channel** command or the Red Hat Network web interface to remove the **Red Hat Enterprise Virtualization Manager (v.3.2 x86\_64)** channels.

```
# rhn-channel --remove --channel=rhel-x86_64-server-6-rhev-3.2
```

5. Run the following command to ensure all packages related to Red Hat Enterprise Virtualization are up to date:

```
# yum update
```

In particular, if you are using the JBoss Application Server from JBoss Enterprise Application Platform 6.0.1, you must run the above command to upgrade to Enterprise Application Platform 6.1.

## Result

Red Hat Enterprise Virtualization Manager has been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.3 features you must also:

- Ensure all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.3.
- Change all of your data centers to use compatibility version 3.3.

[Report a bug](#)

## 5.4. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER 3.2

### 5.4.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2


#### Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.2 is performed using the **rhevms-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.



#### IMPORTANT

Users of Red Hat Enterprise Virtualization 3.0 must migrate to Red Hat Enterprise Virtualization 3.1 before attempting this upgrade.



#### NOTE

In the event that the upgrade fails the **rhevms-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.

#### Procedure 5.8. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

##### 1. Add Red Hat Enterprise Virtualization 3.2 Subscription

Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.2 packages. This procedure assumes that the system is already subscribed to required channels and entitlements to receive Red Hat Enterprise Virtualization 3.1 packages. These must also be available to complete the upgrade process.

##### Certificate-based Red Hat Network

The Red Hat Enterprise Virtualization 3.2 packages are provided by the **rhel-6-server-rhevms-3.2-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The

**yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --enable rhel-6-server-rhev-3.2-rpms
```

### Red Hat Network Classic

The Red Hat Enterprise Virtualization 3.2 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.2 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.2** in Red Hat Network Classic.

```
rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.2
```

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.2 x86\_64)** channel.

## 2. Remove Enterprise Virtualization 3.1 Subscription

Ensure that the system does not use any Red Hat Enterprise Virtualization Manager 3.1 packages by removing the Red Hat Enterprise Virtualization Manager 3.1 channels and entitlements.

### Certificate-based Red Hat Network

Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.1 repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --disablerepo=rhel-6-server-rhev-3.1-rpms
```

### Red Hat Network Classic

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to remove the **Red Hat Enterprise Virtualization Manager (v.3.1 x86\_64)** channels.

```
# rhn-channel --remove --channel=rhel-6-server-rhev-3.1
```

## 3. Update the **rhev-setup** Package

To ensure that you have the most recent version of the **rhev-upgrade** command installed you must update the **rhev-setup** package. Log in as the **root** user and use **yum** to update the **rhev-setup** package.

```
# yum update rhev-setup
```

## 4. Run the **rhev-upgrade** Command

To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-upgrade** command. You must be logged in as the **root** user to run this command.

```
# rhev-upgrade
Loaded plugins: product-id, rhnplugin
Info: RHEV Manager 3.1 to 3.2 upgrade detected
Checking pre-upgrade conditions...(This may take several minutes)
```

5. If the **ipa-server** package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.2 does not support installation on the same machine as Identity Management (IdM).

Error: IPA was found to be installed on this machine. Red Hat Enterprise Virtualization Manager 3.2 does not support installing IPA on the same machine. Please remove ipa packages before you continue.

To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see <https://access.redhat.com/knowledge/articles/233143>.

## Result

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.2 features you must also:

- Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.2.
- Change all of your data centers to use compatibility version 3.2.

[Report a bug](#)

## 5.5. UPGRADING TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER 3.1

### 5.5.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.1


#### Summary

Upgrading Red Hat Enterprise Virtualization Manager to version 3.1 is performed using the **rhevms-upgrade** command. Virtualization hosts, and the virtual machines running upon them, will continue to operate independently while the Manager is being upgraded. Once the Manager upgrade is complete you will be able to upgrade your hosts, if you haven't already, to the latest versions of Red Hat Enterprise Linux and Red Hat Enterprise Virtualization Hypervisor.



#### IMPORTANT

Refer to <https://access.redhat.com/knowledge/articles/269333> for an up to date list of tips and considerations to be taken into account when upgrading to Red Hat Enterprise Virtualization 3.1.



#### IMPORTANT

Users of Red Hat Enterprise Virtualization 2.2 must migrate to Red Hat Enterprise Virtualization 3.0 before attempting this upgrade. For information on migrating from Red Hat Enterprise Virtualization 2.2 to Red Hat Enterprise Virtualization 3.0, refer to <https://access.redhat.com/knowledge/techbriefs/migrating-red-hat-enterprise-virtualization-manager-version-22-30>.



#### NOTE

In the event that the upgrade fails the **rhevms-upgrade** command will attempt to roll your Red Hat Enterprise Virtualization Manager installation back to its previous state. Where this also fails detailed instructions for manually restoring the installation are displayed.



## Procedure 5.9. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

### 1. Red Hat JBoss Enterprise Application Platform 6 Subscription

Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat JBoss Enterprise Application Platform 6 packages. Red Hat JBoss Enterprise Application Platform 6 is a required dependency of Red Hat Enterprise Virtualization Manager 3.1.

#### Certificate-based Red Hat Network

The Red Hat JBoss Enterprise Application Platform 6 packages are provided by the **Red Hat JBoss Enterprise Application Platform** entitlement in certificate-based Red Hat Network.

Use the **subscription-manager** command to ensure that the system is subscribed to the **Red Hat JBoss Enterprise Application Platform** entitlement.

```
# subscription-manager list
```

#### Red Hat Network Classic

The Red Hat JBoss Enterprise Application Platform 6 packages are provided by the **Red Hat JBoss Application Platform (v 6) for 6Server x86\_64** channel, also referred to as **jbappplatform-6-x86\_64-server-6-rpm**, in Red Hat Network Classic. The Channel Entitlement Name for this channel is **Red Hat JBoss Enterprise Application Platform (v 4, zip format)**.

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat JBoss Application Platform (v 6) for 6Server x86\_64** channel.

### 2. Add Red Hat Enterprise Virtualization 3.1 Subscription

Ensure that the system is subscribed to the required channels and entitlements to receive Red Hat Enterprise Virtualization Manager 3.1 packages.

#### Certificate-based Red Hat Network

The Red Hat Enterprise Virtualization 3.1 packages are provided by the **rhel-6-server-rhev-3.1-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **yum-config-manager** command to enable the repository in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --enable rhel-6-server-rhev-3.1-rpms
```

#### Red Hat Network Classic

The Red Hat Enterprise Virtualization 3.1 packages are provided by the **Red Hat Enterprise Virtualization Manager (v.3.1 x86\_64)** channel, also referred to as **rhel-x86\_64-server-6-rhev-3.1** in Red Hat Network Classic.

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to subscribe to the **Red Hat Enterprise Virtualization Manager (v.3.1 x86\_64)** channel.

### 3. Remove Red Hat Enterprise Virtualization 3.0 Subscription

Ensure that the system does not use any Red Hat Enterprise Virtualization Manager 3.0 packages by removing the Red Hat Enterprise Virtualization Manager 3.0 channels and entitlements.

#### Certificate-based Red Hat Network

Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.0

Use the **yum-config-manager** command to disable the Red Hat Enterprise Virtualization 3.0 repositories in your **yum** configuration. The **yum-config-manager** command must be run while logged in as the **root** user.

```
# yum-config-manager --disablerepo=rhel-6-server-rhev-3-rpms
```

```
# yum-config-manager --disablerepo=jb-eap-5-for-rhel-6-server-rpms
```

### Red Hat Network Classic

Use the **rhn-channel** command, or the Red Hat Network Web Interface, to remove the **Red Hat Enterprise Virtualization Manager (v.3.0 x86\_64)** channels.

```
# rhn-channel --remove --channel=rhel-6-server-rhev-3
```

```
# rhn-channel --remove --channel=jbappplatform-5-x86_64-server-6-rpm
```

#### 4. Update the **rhev-setup** Package

To ensure that you have the most recent version of the **rhev-setup** command installed you must update the **rhev-setup** package. Log in as the **root** user and use **yum** to update the **rhev-setup** package.

```
# yum update rhev-setup
```

#### 5. Run the **rhev-upgrade** Command

To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-upgrade** command. You must be logged in as the **root** user to run this command.

```
# rhev-upgrade
Loaded plugins: product-id, rhnplugin
Info: RHEV Manager 3.0 to 3.1 upgrade detected
Checking pre-upgrade conditions...(This may take several minutes)
```

6. If the **ipa-server** package is installed then an error message is displayed. Red Hat Enterprise Virtualization Manager 3.1 does not support installation on the same machine as Identity Management (IdM).

```
Error: IPA was found to be installed on this machine. Red Hat Enterprise Virtualization
Manager 3.1 does not support installing IPA on the same machine. Please remove ipa
packages before you continue.
```

To resolve this issue you must migrate the IdM configuration to another system before re-attempting the upgrade. For further information see <https://access.redhat.com/knowledge/articles/233143>.

7. A list of packages that depend on Red Hat JBoss Enterprise Application Platform 5 is displayed. These packages must be removed to install Red Hat JBoss Enterprise Application Platform 6, required by Red Hat Enterprise Virtualization Manager 3.1.

```
Warning: the following packages will be removed if you proceed with the upgrade:
```

```
* objectweb-asm
```

Would you like to proceed? (yes|no):

You must enter **yes** to proceed with the upgrade, removing the listed packages.

### Result

Your Red Hat Enterprise Virtualization Manager installation has now been upgraded. To take full advantage of all Red Hat Enterprise Virtualization 3.1 features you must also:

- Ensure that all of your virtualization hosts are up to date and running the most recent Red Hat Enterprise Linux packages or Hypervisor images.
- Change all of your clusters to use compatibility version 3.1.
- Change all of your data centers to use compatibility version 3.1.

[Report a bug](#)

## 5.6. POST-UPGRADE TASKS

### 5.6.1. Changing the Cluster Compatibility Version

#### Summary

Red Hat Enterprise Virtualization clusters have a compatibility version. The cluster compatibility version indicates the features of Red Hat Enterprise Virtualization supported by all of the hosts in the cluster. The cluster compatibility is set according to the version of the least capable host operating system in the cluster.



#### NOTE

To change the cluster compatibility version, you must have first updated all the hosts in your cluster to a level that supports your desired compatibility level.

#### Procedure 5.10. Changing the Cluster Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Clusters** tab.
3. Select the cluster to change from the list displayed. If the list of clusters is too long to filter visually then perform a search to locate the desired cluster.
4. Click the **Edit** button.
5. Change the **Compatibility Version** to the desired value.
6. Click **OK** to open the **Change Cluster Compatibility Version** confirmation window.
7. Click **OK** to confirm.

#### Result

You have updated the compatibility version of the cluster. Once you have updated the compatibility version of all clusters in a data center, then you are also able to change the compatibility version of the data center itself.



### WARNING

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center. If you are upgrading the compatibility version from below 3.1 to a higher version, these storage domains will become unusable with versions older than 3.1.

[Report a bug](#)

## 5.6.2. Changing the Data Center Compatibility Version

### Summary

Red Hat Enterprise Virtualization data centers have a compatibility version. The compatibility version indicates the version of Red Hat Enterprise Virtualization that the data center is intended to be compatible with. All clusters in the data center must support the desired compatibility level.



### NOTE

To change the data center compatibility version, you must have first updated all the clusters in your data center to a level that supports your desired compatibility level.

### Procedure 5.11. Changing the Data Center Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Data Centers** tab.
3. Select the data center to change from the list displayed. If the list of data centers is too long to filter visually then perform a search to locate the desired data center.
4. Click the **Edit** button.
5. Change the **Compatibility Version** to the desired value.
6. Click **OK**.

### Result

You have updated the compatibility version of the data center.

**WARNING**

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center. If you are upgrading the compatibility version from below 3.1 to a higher version, these storage domains will become unusable with versions older than 3.1.

[Report a bug](#)

## PART III. INSTALLING HOSTS

## CHAPTER 6. INTRODUCTION TO HOSTS

### 6.1. WORKFLOW PROGRESS - INSTALLING VIRTUALIZATION HOSTS



[Report a bug](#)

### 6.2. INTRODUCTION TO VIRTUALIZATION HOSTS

Red Hat Enterprise Virtualization supports both virtualization hosts which run the Red Hat Enterprise Virtualization Hypervisor, and those which run Red Hat Enterprise Linux. Both types of virtualization host are able to coexist in the same Red Hat Enterprise Virtualization environment.

Prior to installing virtualization hosts you should ensure that:

- all virtualization hosts meet the hardware requirements, and
- you have successfully completed installation of the Red Hat Enterprise Virtualization Manager.

Additionally you may have chosen to install the Red Hat Enterprise Virtualization Manager Reports. This is not mandatory and is not required to commence installing virtualization hosts. Once you have completed the above tasks you are ready to install virtualization hosts.



**IMPORTANT**

It is recommended that you install at least two virtualization hosts and attach them to the Red Hat Enterprise Virtualization environment. Where you attach only one virtualization host you will be unable to access features such as migration which require redundant hosts.



**IMPORTANT**

The Red Hat Enterprise Virtualization Hypervisor is a closed system. Use a Red Hat Enterprise Linux host if additional rpms are required for your environment.

[Report a bug](#)



# CHAPTER 7. RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR HOSTS

## 7.1. RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR INSTALLATION OVERVIEW

Before commencing Hypervisor installation you must be aware that:

- The Red Hat Enterprise Virtualization Hypervisor *must* be installed on a physical server. It must not be installed in a Virtual Machine.
- The installation process will reconfigure the selected storage device and destroy all data. Therefore, ensure that any data to be retained is successfully backed up before proceeding.
- All Hypervisors in an environment must have unique hostnames and IP addresses, in order to avoid network conflicts.
- Instructions for using Network (PXE) Boot to install the Hypervisor are contained in the *Red Hat Enterprise Linux - Installation Guide*, available at [https://access.redhat.com/documentation/en-US/Red\\_Hat\\_Enterprise\\_Linux](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux).
- Red Hat Enterprise Virtualization Hypervisors can use Storage Attached Networks (SANs) and other network storage for storing virtualized guest images. However, a local storage device is required for installing and booting the Hypervisor.



### NOTE

Red Hat Enterprise Virtualization Hypervisor installations can be automated or conducted without interaction. This type of installation is only recommended for advanced users.

[Report a bug](#)

## 7.2. INSTALLING THE RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR DISK IMAGE

### Summary

Before you can set up a Red Hat Enterprise Virtualization Hypervisor, you must download the packages containing the Red Hat Enterprise Virtualization Hypervisor disk image and tools for writing that disk image to USB storage devices or preparing that disk image for deployment via PXE.

### Procedure 7.1. Installing the Red Hat Enterprise Virtualization Hypervisor Disk Image

1. Enable the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86\_64)** repository:

- With RHN Classic:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev
```

- With Subscription Manager, attach a **Red Hat Enterprise Virtualization** entitlement and run the following command:

```
# subscription-manager repos --enable=rhel-6-server-rhev-rpms
```

2. Run the following command to install the `rhev-hypervisor6` package:

```
# yum install rhev-hypervisor6
```

3. Run the following command to install the `livecd-tools` package:

```
# yum install livecd-tools
```

## Result

You have installed the Red Hat Enterprise Virtualization Hypervisor disk image and the `livecd-iso-to-disk` and `livecd-iso-to-pxeboot` utilities. By default, the Red Hat Enterprise Virtualization Hypervisor disk image is located in the `/usr/share/rhev-hypervisor/` directory.



## NOTE

Red Hat Enterprise Linux 6.2 and later allows more than one version of the ISO image to be installed at one time. As such, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso` is now a symbolic link to a uniquely-named version of the Hypervisor ISO image, such as `/usr/share/rhev-hypervisor/rhev-hypervisor-6.4-20130321.0.el6ev.iso`. Different versions of the image can now be installed alongside each other, allowing administrators to run and maintain a cluster on a previous version of the Hypervisor while upgrading another cluster for testing. Additionally, the symbolic link `/usr/share/rhev-hypervisor/rhev-latest-6.iso`, is created. This links also targets the most recently installed version of the Red Hat Enterprise Virtualization ISO image.

[Report a bug](#)

## 7.3. PREPARING INSTALLATION MEDIA

### 7.3.1. Preparing a USB Storage Device

You can write the Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device such as a flash drive or external hard drive. You can then use that USB device to start the machine on which the Red Hat Enterprise Virtualization Hypervisor will be installed and install the Red Hat Enterprise Virtualization Hypervisor operating system.



## NOTE

Not all systems support booting from a USB storage device. Ensure the BIOS on the system on which you will install the Red Hat Enterprise Virtualization Hypervisor supports this feature.

[Report a bug](#)

### 7.3.2. Preparing USB Installation Media Using `livecd-iso-to-disk`

## Summary

You can use the **livecd-iso-to-disk** utility included in the `livecd-tools` package to write a hypervisor or other disk image to a USB storage device. You can then use that USB storage device to start systems that support booting via USB and install the Red Hat Enterprise Virtualization Hypervisor.

The basic syntax for the **livecd-iso-to-disk** utility is as follows:

```
# livecd-iso-to-disk [image] [device]
```

The `[device]` parameter is the path to the USB storage device on which to write the disk image. The `[image]` parameter is the path and file name of the disk image to write to the USB storage device. By default, the Red Hat Enterprise Virtualization Hypervisor disk image is located at **/usr/share/rhev-hypervisor/rhev-hypervisor.iso** on the machine on which the Red Hat Enterprise Virtualization Manager is installed. The **livecd-iso-to-disk** utility requires devices to be formatted with the **FAT** or **EXT3** file system.



## NOTE

USB storage devices are sometimes formatted without a partition table. In this case, use a generic identifier for the storage device such as **/dev/sdb**. When a USB storage device is formatted with a partition table, use the path name to the device, such as **/dev/sdb1**.

## Procedure 7.2. Preparing USB Installation Media Using livecd-iso-to-disk

1. Run the following command to ensure you have the latest version of the Red Hat Enterprise Virtualization Hypervisor disk image:

```
# yum update rhv-hypervisor6
```

2. Use the **livecd-iso-to-disk** utility to write the disk image to a USB storage device.

### Example 7.1. Use of livecd-iso-to-disk

This example demonstrates the use of **livecd-iso-to-disk** to write a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device named **/dev/sdc** and make that USB storage device bootable.

```
# livecd-iso-to-disk --format --reset-mbr /usr/share/rhev-hypervisor/rhev-hypervisor.iso
/dev/sdc
Verifying image...
/usr/share/rhev-hypervisor/rhev-hypervisor.iso:  eccc12a0530b9f22e5ba62b848922309
Fragment sums: 8688f5473e9c176a73f7a37499358557e6c397c9ce2dafb5eca5498fb586
Fragment count: 20
Press [Esc] to abort check.
Checking: 100.0%
```

The media check is complete, the result is: PASS.

It is OK to use this media.

```
WARNING: THIS WILL DESTROY ANY DATA ON /dev/sdc!!!
Press Enter to continue or ctrl-c to abort
```

```
/dev/sdc: 2 bytes were erased at offset 0x000001fe (dos): 55 aa
Waiting for devices to settle...
```

```
mke2fs 1.42.7 (21-Jan-2013)
Filesystem label=LIVE
OS type: Linux
Block size=4096 (log=2)
Fragment size=4096 (log=2)
Stride=0 blocks, Stripe width=0 blocks
488640 inodes, 1953280 blocks
97664 blocks (5.00%) reserved for the super user
First data block=0
Maximum filesystem blocks=2000683008
60 block groups
32768 blocks per group, 32768 fragments per group
8144 inodes per group
Superblock backups stored on blocks:
    32768, 98304, 163840, 229376, 294912, 819200, 884736, 1605632

Allocating group tables: done
Writing inode tables: done
Creating journal (32768 blocks): done
Writing superblocks and filesystem accounting information: done

Copying live image to target device.
squashfs.img
 163360768 100% 184.33MB/s  0:00:00 (xfer#1, to-check=0/1)

sent 163380785 bytes received 31 bytes 108920544.00 bytes/sec
total size is 163360768 speedup is 1.00
osmin.img
 4096 100% 0.00kB/s  0:00:00 (xfer#1, to-check=0/1)

sent 4169 bytes received 31 bytes 8400.00 bytes/sec
total size is 4096 speedup is 0.98
Updating boot config file
Installing boot loader
/media/tgttmp.q6aZdS/syslinux is device /dev/sdc
Target device is now set up with a Live image!
```

## Result

You have written a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device. You can now use that USB storage device to start a system and install the Red Hat Enterprise Virtualization Hypervisor operating system.

[Report a bug](#)

### 7.3.3. Preparing USB Installation Media Using `dd`

The `dd` utility can also be used to write a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device. The `dd` utility is available from the `coreutils` package, and versions of the `dd` utility are available on a wide variety of Linux and Unix operating systems. Windows users can obtain the `dd` utility by installing Red Hat Cygwin, a free Linux-like environment for Windows.

The basic syntax for the `dd` utility is as follows:

```
# dd if=[image] of=[device]
```

The `[device]` parameter is the path to the USB storage device on which the disk image will be written. The `[image]` parameter is the path and file name of the disk image to write to the USB storage device. By default, the Red Hat Enterprise Virtualization Hypervisor disk image is located at `/usr/share/rhev-hypervisor/rhev-hypervisor.iso` on the machine on which the `rhev-hypervisor6` package is installed. The `dd` command does not make assumptions as to the format of the device because it performs a low-level copy of the raw data in the selected image.

[Report a bug](#)

### 7.3.4. Preparing USB Installation Media Using `dd` on Linux Systems

#### Summary

You can use the `dd` utility to write a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device.

#### Procedure 7.3. Preparing USB Installation Media using `dd` on Linux Systems

1. Run the following command to ensure you have the latest version of the Red Hat Enterprise Virtualization Hypervisor disk image:

```
# yum update rhv-hypervisor6
```

2. Use the `dd` utility to write the disk image to a USB storage device.

#### Example 7.2. Use of `dd`

This example uses a USB storage device named `/dev/sdc`.

```
# dd if=/usr/share/rhev-hypervisor/rhev-hypervisor.iso of=/dev/sdc
243712+0 records in
243712+0 records out
124780544 bytes (125 MB) copied, 56.3009 s, 2.2 MB/s
```



#### WARNING

The `dd` utility will overwrite all data on the device specified by the `of` parameter. Ensure you have specified the correct device and that the device contains no valuable data before using the `dd` utility.

#### Result

You have written a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device.

[Report a bug](#)


### 7.3.5. Preparing USB Installation Media Using `dd` on Windows Systems

## Summary

You can use the **dd** utility to write a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device. To use this utility in Windows, you must download and install Red Hat Cygwin.

### Procedure 7.4. Preparing USB Installation Media using dd on Windows Systems

1. Open <http://www.redhat.com/services/custom/cygwin/> in a web browser and click **32-bit Cygwin** to download the 32-bit version of Red Hat Cygwin, or **64-bit Cygwin** to download the 64-bit version of Red Hat Cygwin.
2. Run the downloaded executable as a user with administrator privileges to open the Red Hat Cygwin installation program.
3. Follow the prompts to install Red Hat Cygwin. The Coreutils package in the Base package group provides the **dd** utility. This package is automatically selected for installation.
4. Copy the **rhev-hypervisor.iso** file downloaded from the Red Hat Network to **C:\rhev-hypervisor.iso**.
5. Run the Red Hat Cygwin application from the desktop as a user with administrative privileges.



#### IMPORTANT

On the Windows 7 and Windows Server 2008, you must right-click the Red Hat Cygwin icon and select the **Run as Administrator** option to ensure the application runs with the correct permissions.

6. In the terminal, run the following command to view the drives and partitions currently visible to the system:

```
$ cat /proc/partitions
```

#### Example 7.3. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor #blocks name
 8     0 15728640 sda
 8     1  102400 sda1
 8     2 15624192 sda2
```

7. Attach the USB storage device to which the Red Hat Enterprise Virtualization Hypervisor disk image will be written to the system. Run the **cat /proc/partitions** command again and compare the output to that of the previous output. A new entry will appear that designates the USB storage device.

#### Example 7.4. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor #blocks name
 8     0 15728640 sda
```

```
8 1 102400 sda1
8 2 15624192 sda2
8 16 524288 sdb
```

- Use the **dd** utility to write the **rhev-hypervisor.iso** file to the USB storage device. The following example uses a USB storage device named **/dev/sdb**. Replace *sdb* with the correct device name for the USB storage device to be used.


#### Example 7.5. Use of dd Utility Under Red Hat Cygwin

```
Administrator@test /
$ dd if=/cygdrive/c/rhev-hypervisor.iso of=/dev/sdb& pid=$!
```



#### WARNING

The **dd** utility will overwrite all data on the device specified by the **of** parameter. Ensure you have specified the correct device and that the device contains no valuable data before using the **dd** utility.



#### NOTE

Writing disk images to USB storage devices with the version of the **dd** utility included with Red Hat Cygwin can take significantly longer than the equivalent on other platforms. You can run the following command to view the progress of the operation:

```
$ kill -USR1 $pid
```

#### Result

You have written a Red Hat Enterprise Virtualization Hypervisor disk image to a USB storage device.

[Report a bug](#)

### 7.3.6. Preparing Optical Hypervisor Installation Media

#### Summary

You can write a Red Hat Enterprise Virtualization Hypervisor disk image to a CD-ROM or DVD with the **wodim** utility. The **wodim** utility is provided by the **wodim** package.

#### Procedure 7.5. Preparing Optical Hypervisor Installation Media

- Run the following command to install the **wodim** package and dependencies:

```
# yum install wodim
```

2. Insert a blank CD-ROM or DVD into your CD or DVD writer.
3. Run the following command to write the Red Hat Enterprise Virtualization Hypervisor disk image to the disc:

```
wodim dev=[device] [image]
```

#### Example 7.6. Use of the wodim Utility

This example uses the first CD-RW (**/dev/cdrw**) device available and the default hypervisor image location.

```
# wodim dev=/dev/cdrw /usr/share/rhev-hypervisor/rhev-hypervisor.iso
```



#### IMPORTANT

The Hypervisor uses a program (**isomd5sum**) to verify the integrity of the installation media every time the hypervisor is booted. If media errors are reported in the boot sequence you have a bad CD-ROM. Follow the procedure above to create a new CD-ROM or DVD.

#### Result

You have written a Red Hat Enterprise Virtualization Hypervisor disk image to a CD-ROM or DVD.

[Report a bug](#)

## 7.4. INSTALLATION

### 7.4.1. Booting the Hypervisor from USB Installation Media

#### Summary

Booting a hypervisor from a USB storage device is similar to booting other live USB operating systems. Follow this procedure to boot a machine using USB installation media.

#### Procedure 7.6. Booting the Hypervisor from USB Installation Media

1. Enter the BIOS menu to enable USB storage device booting if not already enabled.
  - a. Enable USB booting if this feature is disabled.
  - b. Set booting USB storage devices to be first boot device.
  - c. Shut down the system.
2. Insert the USB storage device that contains the hypervisor boot image.
3. Restart the system.

#### Result

The hypervisor boot process commences automatically.



[Report a bug](#)

## 7.4.2. Booting the Hypervisor from Optical Installation Media

### Summary

Booting the Hypervisor from optical installation media requires the system to have a correctly defined BIOS boot configuration.

### Procedure 7.7. Booting the Hypervisor from Optical Installation Media

1. Ensure that the system's BIOS is configured to boot from the CD-ROM or DVD-ROM drive first. For many systems this the default.



#### NOTE

Refer to your manufacturer's manuals for further information on modifying the system's BIOS boot configuration.

2. Insert the Hypervisor CD-ROM in the CD-ROM or DVD-ROM drive.
3. Reboot the system.

### Result

The Hypervisor boot screen will be displayed.

[Report a bug](#)

## 7.4.3. Starting the Installation Program

### Summary

When you start a system using the prepared boot media, the first screen that displays is the boot menu. From here, you can start the installation program for installing the hypervisor.

### Procedure 7.8. Starting the Installation Program

1. From the boot splash screen, press any key to open the boot menu.




Figure 7.1. The boot splash screen

2. From the boot menu, use the directional keys to select **Install or Upgrade, Install (Basic Video)**, or **Install or Upgrade with Serial Console**.

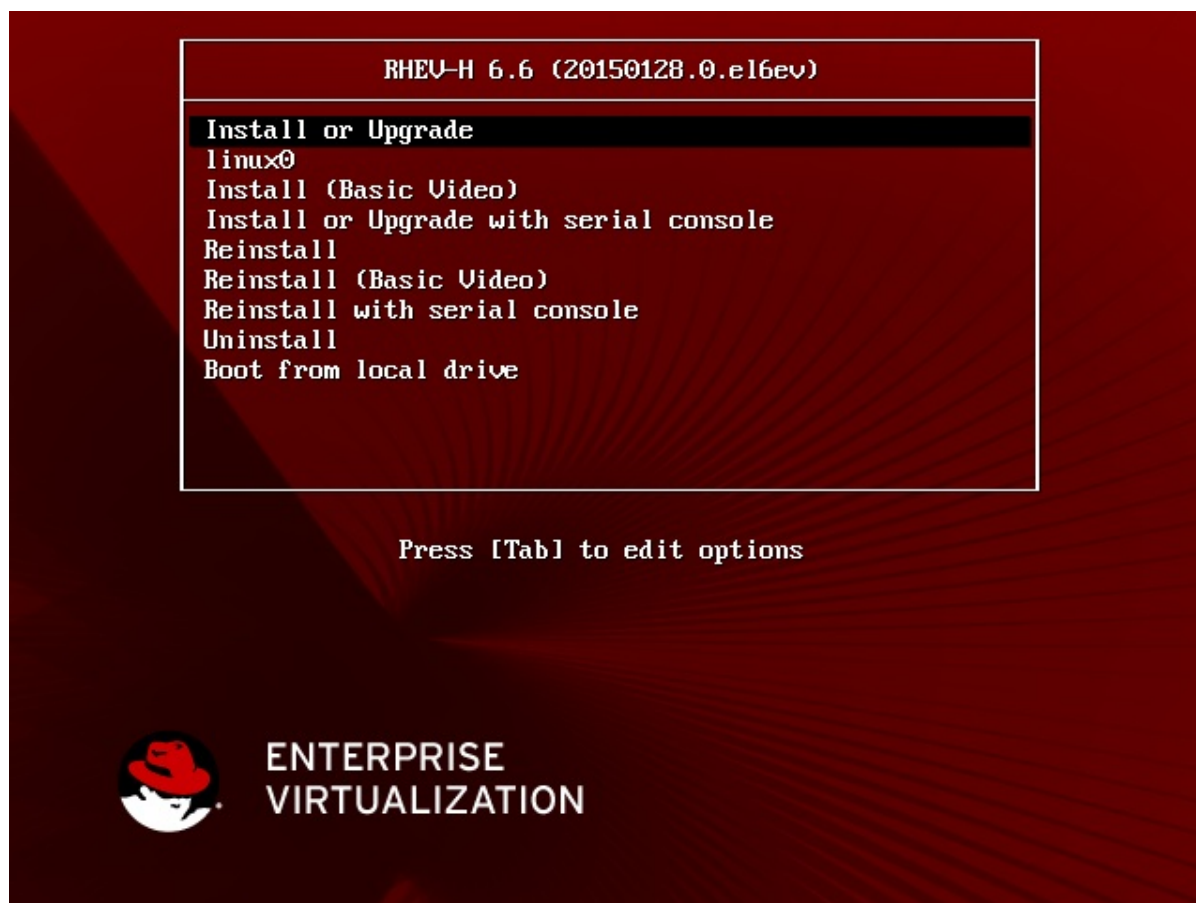


Figure 7.2. The boot menu

The full list of options in the boot menu is as follows:

#### **Install or Upgrade**

Install or upgrade the hypervisor.

#### **Install (Basic Video)**

Install or upgrade the Hypervisor in basic video mode.

#### **Install or Upgrade with Serial Console**

Install or upgrade the hypervisor while redirecting the console to a serial device attached to `/dev/ttyS0`.

#### **Reinstall**

Reinstall the hypervisor.

#### **Reinstall (Basic Video)**

Reinstall the hypervisor in basic video mode.

#### **Reinstall with Serial Console**

Reinstall the hypervisor while redirecting the console to a serial device attached to `/dev/ttyS0`.

#### **Uninstall**

Uninstall the hypervisor.

## Boot from Local Drive

Boot the operating system installed on the first local drive.

3. Press the **Enter** key.



### NOTE

From the boot menu, you can also press the **Tab** key to edit the kernel parameters. Kernel parameters must be separated by a space, and once you have entered the preferred kernel parameters, you can boot the system using those kernel parameters by pressing the **Enter** key. To clear any changes you have made to the kernel parameters and return to the boot menu, press the **Esc** key.

### Result

You have started the hypervisor installation program.

[Report a bug](#)

## 7.4.4. Hypervisor Menu Actions

- The directional keys (**Up**, **Down**, **Left**, **Right**) are used to select different controls on the screen. Alternatively the **Tab** key cycles through the controls on the screen which are enabled.
- Text fields are represented by a series of underscores (`_`). To enter data in a text field select it and begin entering data.
- Buttons are represented by labels which are enclosed within a pair of angle brackets (`<` and `>`). To activate a button ensure it is selected and press **Enter** or **Space**.
- Boolean options are represented by an asterisk (\*) or a space character enclosed within a pair of square brackets (`[` and `]`). When the value contained within the brackets is an asterisk then the option is set, otherwise it is not. To toggle a Boolean option on or off press **Space** while it is selected.

[Report a bug](#)

## 7.4.5. Installing the Hypervisor

### Summary

There are two methods for installing Red Hat Enterprise Virtualization Hypervisors:

- Interactive installation.
- Unattended installation.

This section outlines the procedure for installing a Hypervisor interactively.

### Procedure 7.9. Installing the Hypervisor Interactively

1. Use the prepared boot media to boot the machine on which the Hypervisor is to be installed.
2. Select **Install Hypervisor** and press **Enter** to begin the installation process.

3. The first screen that appears allows you to configure the appropriate keyboard layout for your locale. Use the arrow keys to highlight the appropriate option and press **Enter** to save your selection.

#### Example 7.7. Keyboard Layout Configuration

```
Keyboard Layout Selection

Available Keyboard Layouts
Swiss German (latin1)
Turkish
U.S. English
U.S. International
...

(Hit enter to select a layout)

<Quit>  <Back>  <Continue>
```

4. The installation script automatically detects all disks attached to the system. This information is used to assist with selection of the boot and installation disks that the Hypervisor will use. Each entry displayed on these screens indicates the **Location**, **Device Name**, and **Size** of the disks.

##### a. Boot Disk

The first disk selection screen is used to select the disk from which the Hypervisor will boot. The Hypervisor's boot loader will be installed to the Master Boot Record (MBR) of the disk that is selected on this screen. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which to choose the boot device. Alternatively, you can manually select a device by specifying a block device name using the **Other Device** option.



#### IMPORTANT

The selected disk must be identified as a boot device and appear in the boot order either in the system's BIOS or in a pre-existing boot loader.

##### ■ Automatically Detected Device Selection

- i. Select the entry for the disk the Hypervisor is to boot from in the list and press **Enter**.
- ii. Select the disk and press **Enter**. This action saves the boot device selection and starts the next step of installation.

##### ■ Manual Device Selection

- i. Select **Other device** and press **Enter**.
- ii. When prompted to **Please select the disk to use for booting RHEV-H**, enter the name of the block device from which the Hypervisor should boot.

#### Example 7.8. Other Device Selection



Please select the disk to use for booting RHEV-H  
/dev/sda

- iii. Press **Enter**. This action saves the boot device selection and starts the next step of installation.
- b. The disk or disks selected for installation will be those to which the Hypervisor itself is installed. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which installation devices are chosen.



#### WARNING

All data on the selected storage devices will be destroyed.

- i. Select each disk on which the Hypervisor is to be installed and press **Space** to toggle it to enabled. Where other devices are to be used for installation, either solely or in addition to those which are listed automatically, use **Other Device**.
- ii. Select the **Continue** button and press **Enter** to continue.
- iii. Where the **Other Device** option was specified, a further prompt will appear. Enter the name of each additional block device to use for Hypervisor installation, separated by a comma. Once all required disks have been selected, select the **<Continue>** button and press **Enter**.

#### Example 7.9. Other Device Selection

Please enter one or more disks to use for installing RHEV-H. Multiple devices can be separated by comma.  
Device path: /dev/mmcbk0,/dev/mmcbk1\_\_\_\_\_

Once the installation disks have been selected, the next stage of the installation starts.

5. The next screen allows you to configure storage for the Hypervisor.
  1. Select or clear the **Fill disk with Data partition** check box. Clearing this text box displays a field showing the remaining space on the drive and allows you to specify the amount of space to be allocated to data storage.
  2. Enter the preferred values for **Swap, Config, and Logging**.
  3. If you selected the **Fill disk with Data partition** check box, the **Data** field is automatically set to **0**. If the check box was cleared, you can enter a whole number up to the value of the **Remaining Space** field. Entering a value of **-1** fills all remaining space.

- The Hypervisor requires a password be set to protect local console access to the **admin** user. The installation script prompts you to enter the preferred password in both the **Password** and **Confirm Password** fields.

Use a strong password. Strong passwords comprise a mix of uppercase, lowercase, numeric, and punctuation characters. They are six or more characters long and do not contain dictionary words.

Once a strong password has been entered, select **<Install>** and press **Enter** to install the Hypervisor on the selected disks.


## Result

Once installation is complete, the message **RHEV Hypervisor Installation Finished Successfully** will be displayed. Select the **<Reboot>** button and press **Enter** to reboot the system.



### NOTE

Remove the boot media and change the boot device order to prevent the installation sequence restarting after the system reboots.



### NOTE

Red Hat Enterprise Virtualization Hypervisors are able to use Storage Area Networks (SANs) and other network storage for storing virtualized guest images. Hypervisors can be installed on SANs, provided that the Host Bus Adapter (HBA) permits configuration as a boot device in BIOS.



### NOTE

Hypervisors are able to use multipath devices for installation. Multipath is often used for SANs or other networked storage. Multipath is enabled by default at install time. Any block device which responds to **scsi\_id** functions with multipath. Devices where this is not the case include USB storage and some older ATA disks.

[Report a bug](#)

## 7.5. CONFIGURATION

### 7.5.1. Logging Into the Hypervisor

#### Summary

You can log into the hypervisor console locally to configure the hypervisor.

#### Procedure 7.10. Logging Into the Hypervisor

- Start the machine on which the Red Hat Enterprise Virtualization Hypervisor operating system is installed.
- Enter the user name **admin** and press **Enter**.
- Enter the password you set during installation and press **Enter**.

## Result

You have successfully logged into the hypervisor console as the **admin** user.

[Report a bug](#)

## 7.5.2. The Status Screen

The **Status** screen provides an overview of the state of the Hypervisor such as the current status of networking, the location in which logs and reports are stored, and the number of virtual machines that are active on that hypervisor. The **Status** screen also provides the following buttons for viewing further details regarding the Hypervisor and for changing the state of the Hypervisor:

- **<View Host Key>**: Displays the RSA host key fingerprint and host key of the Hypervisor.
- **<View CPU Details>**: Displays details on the CPU used by the Hypervisor such as the CPU name and type.
- **<Lock>**: Locks the Hypervisor. The user name and password must be entered to unlock the Hypervisor.
- **<Log Off>**: Logs off the current user.
- **<Restart>**: Restarts the Hypervisor.
- **<Power Off>**: Turns the Hypervisor off.

[Report a bug](#)

## 7.5.3. The Network Screen

### 7.5.3.1. The Network Screen

The **Network** screen is used to configure the host name of the hypervisor and the DNS servers, NTP servers and network interfaces that the hypervisor will use. The **Network** screen also provides a number of buttons for testing and configuring network interfaces:

- **<Ping>**: Allows you to ping a given IP address by specifying the address to ping and the number of times to ping that address.
- **<Create Bond>**: Allows you to create bonds between network interfaces.

[Report a bug](#)

### 7.5.3.2. Configuring the Host Name

#### Summary

You can change the host name used to identify the hypervisor.

#### Procedure 7.11. Configuring the Host Name

1. Select the **Hostname** field on the **Network** screen and enter the new host name.
2. Select **<Save>** and press **Enter** to save the changes.



## Result

You have changed the host name used to identify the hypervisor.

[Report a bug](#)

### 7.5.3.3. Configuring Domain Name Servers

#### Summary

You can specify up to two domain name servers that the hypervisor will use to resolve network addresses.

#### Procedure 7.12. Configuring Domain Name Servers

1. To set or change the primary DNS server, select the **DNS Server 1** field and enter the IP address of the new primary DNS server.
2. To set or change the secondary DNS server, select the **DNS Server 2** field and enter the IP address of the new secondary DNS server.
3. Select **<Save>** and press **Enter** to save the changes.

## Result

You have specified the primary and secondary domain name servers that the hypervisor will use to resolve network addresses.

[Report a bug](#)

### 7.5.3.4. Configuring Network Time Protocol Servers

#### Summary

You can specify up to two network time protocol servers that the hypervisor will use to synchronize its system clock.



#### IMPORTANT

You must specify the same time servers as the Red Hat Enterprise Virtualization Manager to ensure all system clocks throughout the Red Hat Enterprise Virtualization environment are synchronized.

#### Procedure 7.13. Configuring Network Time Protocol Servers

1. To set or change the primary NTP server, select the **NTP Server 1** field and enter the IP address or host name of the new primary NTP server.
2. To set or change the secondary NTP server, select the **NTP Server 2** field and enter the IP address or host name of the new secondary NTP server.
3. Select **<Save>** and press **Enter** to save changes to the NTP configuration.

## Result

You have specified the primary and secondary NTP servers that the hypervisor will use to synchronize its system clock.

[Report a bug](#)

### 7.5.3.5. Configuring Network Interfaces

#### Summary

After you have installed the Red Hat Enterprise Virtualization Hypervisor operating system, all network interface cards attached to the hypervisor are initially in an unconfigured state. You must configure at least one network interface to connect the hypervisor with the Red Hat Enterprise Virtualization Manager.

#### Procedure 7.14. Configuring Network Interfaces

1. Select a network interfaces from the list beneath **Available System NICs** and press **Enter** to configure that network interface.



#### NOTE

To identify the physical network interface card associated with the selected network interface, select **<Flash Lights to Identify>** and press **Enter**.

2. Configure a dynamic or static IP address:
  - **Configuring a Dynamic IP Address**  
Select **DHCP** under **IPv4 Settings** and press the space bar to enable this option.
  - **Configuring a Static IP Address**
    - Select **Static** under **IPv4 Settings** and press the space bar to enable this option.
    - Specify the **IP Address**, **Netmask**, and **Gateway** that the hypervisor will use.

#### Example 7.10. Static IPv4 Networking Configuration

```
IPv4 Settings
( ) Disabled  ( ) DHCP  (*) Static
IP Address: 192.168.122.100_ Netmask: 255.255.255.0____
Gateway 192.168.1.1_____
```



#### NOTE

The Red Hat Enterprise Virtualization Manager does not currently support IPv6 networking. IPv6 networking must remain set to **Disabled**.

3. Enter a VLAN identifier in the **VLAN ID** field to configure a VLAN for the device.
4. Select the **Use Bridge** option and press the space bar to enable this option.
5. Select the **<Save>** button and press **Enter** to save the network configuration.

#### Result

The progress of configuration is displayed on screen. When configuration is complete, press the **Enter** key to close the progress window and return to the **Network** screen. The network interface is now listed as **Configured**.

[Report a bug](#)

## 7.5.4. The Security Screen

### Summary

You can configure security-related options for the hypervisor such as SSH password authentication, AES-NI encryption, and the password of the **admin** user.

### Procedure 7.15. Configuring Security

1. Select the **Enable SSH password authentication** option and press the space bar to enable SSH authentication.
2. Select the **Disable AES-NI** option and press the space bar to disable the use of AES-NI for encryption.
3. Optionally, enter the number of bytes by which to pad blocks in AES-NI encryption if AES-NI encryption is enabled.
4. Enter a new password for the **admin** user in the **Password** field and **Confirm Password** to change the password used to log into the hypervisor console.
5. Select **<Save>** and press **Enter**.

### Result

You have updated the security-related options for the hypervisor.

[Report a bug](#)

## 7.5.5. The Keyboard Screen

### Summary

The **Keyboard** screen allows you to configure the keyboard layout used inside the hypervisor console.

### Procedure 7.16. Configuring the Hypervisor Keyboard Layout

1. Select a keyboard layout from the list provided.

Keyboard Layout Selection

Choose the Keyboard Layout you would like to apply to this system.

Current Active Keyboard Layout: U.S. English

Available Keyboard Layouts

Swiss German (latin1)

Turkish

U.S. English

U.S. International

Ukrainian

```

| ...
| <Save>

```

2. Select **Save** and press **Enter** to save the selection.

### Result

You have successfully configured the keyboard layout.

[Report a bug](#)

## 7.5.6. The SNMP Screen

### Summary

The **SNMP** screen allows you to enable and configure a password for simple network management protocol.

```

| Enable SNMP  [ ]
|
| SNMP Password
| Password:  _____
| Confirm Password: _____
|
| <Save>  <Reset>

```

### Procedure 7.17. Configuring Simple Network Management Protocol

1. Select the **Enable SNMP** option and press the space bar to enable SNMP.
2. Enter a password in the **Password** and **Confirm Password** fields.
3. Select **<Save>** and press **Enter**.

### Result

You have enabled SNMP and configured a password that the hypervisor will use in SNMP communication.

[Report a bug](#)

## 7.5.7. The CIM Screen

### Summary

The **CIM** screen allows you to configure a common information model for attaching the hypervisor to a pre-existing CIM management infrastructure and monitor virtual machines that are running on the hypervisor.

### Procedure 7.18. Configuring Hypervisor Common Information Model

1. Select the **Enable CIM** option and press the space bar to enable CIM.

```

| Enable CIM  [ ]

```

2. Enter a password in the **Password** field and **Confirm Password** field.
3. Select **Save** and press **Enter**.

### Result

You have configured the Hypervisor to accept CIM connections authenticated using a password. Use this password when adding the Hypervisor to your common information model object manager.

[Report a bug](#)

## 7.5.8. The Logging Screen

### Summary

The **Logging** screen allows you to configure logging-related options such as a daemon for automatically exporting log files generated by the hypervisor to a remote server.

### Procedure 7.19. Configuring Logging

1. In the **Logrotate Max Log Size** field, enter the maximum size in kilobytes that log files can reach before they are rotated by **logrotate**. The default value is **1024**.
2. Optionally, configure **rsyslog** to transmit log files to a remote **syslog** daemon:
  - a. Enter the remote **rsyslog** server address in the **Server Address** field.
  - b. Enter the remote **rsyslog** server port in the **Server Port** field. The default port is **514**.
3. Optionally, configure **netconsole** to transmit kernel messages to a remote destination:
  - a. Enter the **Server Address**.
  - b. Enter the **Server Port**. The default port is **6666**.
4. Select **<Save>** and press **Enter**.

### Result

You have configured logging for the hypervisor.

[Report a bug](#)

## 7.5.9. The Kdump Screen

### Summary

The **Kdump** screen allows you to specify a location in which kernel dumps will be stored in the event of a system failure. There are four options - **Disable**, which disables kernel dumping, **Local**, which stores kernel dumps on the local system, and **SSH** and **NFS**, which allow you to export kernel dumps to a remote location.

### Procedure 7.20. Configuring Kernel Dumps

1. Select an option for storing kernel dumps:
  - **Local**

- a. Select the **Local** option and press the space bar to store kernel dumps on the local system.
  - **SSH**
    - a. Select the **SSH** option and press the space bar to export kernel dumps via SSH.
    - b. Enter the location in which kernel dumps will be stored in the **SSH Location (root@example.com)** field.
  - **NFS**
    - a. Select the **NFS** option and press the space bar to export kernel dumps to an NFS share.
    - b. Enter the location in which kernel dumps will be stored in the **NFS Location (example.com:/var/crash)** field.
2. Select **<Save>** and press **Enter**.

### Result

You have configured a location in which kernel dumps will be stored in the event of a system failure.

[Report a bug](#)

## 7.5.10. The Remote Storage Screen

### Summary

You can use the **Remote Storage** screen to specify a remote iSCSI initiator or NFS share to use as storage.

### Procedure 7.21. Configuring Remote Storage

1. Enter an initiator name in the **iSCSI Initiator Name** field or the path to the NFS share in the **NFSv4 Domain (example.redhat.com)** field.

#### Example 7.11. iSCSI Initiator Name

```
iSCSI Initiator Name:  
iqn.1994-05.com.redhat:5189835eeb40_____
```

#### Example 7.12. NFS Path

```
NFSv4 Domain (example.redhat.com):  
example.redhat.com_____
```

2. Select **<Save>** and press **Enter**.

### Result

You have configured remote storage.

[Report a bug](#)

### 7.5.11. The Diagnostics Screen

The **Diagnostics** screen allows you to select one of the diagnostic tools from the following list:

- **multipath -ll**: Shows the current multipath topology from all available information.
- **fdisk -l**: Lists the partition tables.
- **parted -s -l**: Lists partition layout on all block devices.
- **lsblk**: Lists information on all block devices.

[Report a bug](#)

### 7.5.12. The Performance Screen

The **Performance** screen allows you to select and apply a **tuned** profile to your system from the following list. The **virtual-host** profile is used by default.

Table 7.1. Tuned Profiles available in Red Hat Enterprise Virtualization


Tuned Profile	Description
<b>None</b>	The system is disabled from using any <b>tuned</b> profile.
<b>virtual-host</b>	Based on the <b>enterprise-storage</b> profile, <b>virtual-host</b> decreases the swappiness of virtual memory and enables more aggressive writeback of dirty pages.
<b>virtual-guest</b>	A profile optimized for virtual machines.
<b>throughput-performance</b>	A server profile for typical throughput performance tuning.
<b>spindown-disk</b>	A strong power-saving profile directed at machines with classic hard disks.
<b>server-powersave</b>	A power-saving profile directed at server systems.
<b>latency-performance</b>	A server profile for typical latency performance tuning.
<b>laptop-battery-powersave</b>	A high-impact power-saving profile directed at laptops running on battery.
<b>laptop-ac-powersave</b>	A medium-impact power-saving profile directed at laptops running on AC.
<b>enterprise-storage</b>	A server profile to improve throughput performance for enterprise-sized server configurations.
<b>desktop-powersave</b>	A power-saving profile directed at desktop systems.

Tuned Profile	Description
<b>default</b>	The default power-saving profile. This is the most basic power-saving profile. It only enables the disk and CPU plug-ins.

[Report a bug](#)

### 7.5.13. The RHEV-M Screen

You can attach the Hypervisor to the Red Hat Enterprise Virtualization Manager immediately if the address of the Manager is available. If the Manager has not yet been installed, you must instead set a password. This allows the Hypervisor to be added from the Administration Portal once the Manager has been installed. Both modes of configuration are supported from the **RHEV-M** screen in the Hypervisor user interface. However, adding the Hypervisor from the Administration Portal is the recommended option.



#### IMPORTANT

Setting a password on the **RHEV-M** configuration screen sets the **root** password on the Hypervisor and enables SSH password authentication. Once the Hypervisor has successfully been added to the Manager, disabling SSH password authentication is recommended.



#### IMPORTANT

If you are configuring the Hypervisor to use a bond or bridge device, add it manually from the Red Hat Enterprise Virtualization Manager instead of registering it with the Manager during setup to avoid unexpected errors.

#### Procedure 7.22. Configuring a Hypervisor Management Server

- ○ Configure the Hypervisor Management Server using the address of the Manager.
  - a. Enter the IP address or fully qualified domain name of the Manager in the **Management Server** field.
  - b. Enter the management server port in the **Management Server Port** field. The default value is **443**. If a different port was selected during Red Hat Enterprise Virtualization Manager installation, specify it here, replacing the default value.
  - c. Leave the **Password** and **Confirm Password** fields blank. These fields are not required if the address of the management server is known.
  - d. Select **<Save & Register>** and press **Enter**.
  - e. In the **RHEV-M Fingerprint** screen, review the SSL fingerprint retrieved from the Manager, select **<Accept>**, and press **Enter**. The **Certificate Status** in the **RHEV-M** screen changes from **N/A** to **Verified**.
- Configure the Hypervisor Management Server using a password.
  - a. Enter a password in the **Password** field. Although the Hypervisor will accept a weak



password, it is recommended that you use a strong password. Strong passwords contain a mix of uppercase, lowercase, numeric and punctuation characters. They are six or more characters long and do not contain dictionary words.

- b. Re-enter the password in the **Confirm Password** field.
- c. Leave the **Management Server** and **Management Server Port** fields blank. As long as a password is set, allowing the Hypervisor to be added to the Manager later, these fields are not required.
- d. Select **<Save & Register>** and press **Enter**.

[Report a bug](#)

### 7.5.14. The Plugins Screen

The **Plugins** screen provides an overview of the installed plug-ins and allows you to view package differences if you have used the **edit-node** tool to update or add new packages. The **Plugins** screen also provides the following buttons:


- **<RPM Diff>**: Allows you to view RPM differences.
- **<SRPM Diff>**: Allows you to view SRPM differences.
- **<File Diff>**: Allows you to view file differences.

[Report a bug](#)

### 7.5.15. The RHN Registration Screen

#### Summary

Guests running on the Hypervisor may need to consume Red Hat Enterprise Linux virtualization entitlements. In this case, the Hypervisor must be registered to Red Hat Network, a Satellite server, or Subscription Asset Manager. The Hypervisor can also connect to these services via a proxy server.



#### NOTE

You do not need to register the hypervisor with the Red Hat Network to receive updates to the hypervisor image itself; new versions of the hypervisor image are made available through the Red Hat Enterprise Virtualization Manager.

#### Procedure 7.23. Registering the Hypervisor with the Red Hat Network

1. Enter your Red Hat Network user name in the **Login** field.
2. Enter your Red Hat Network password in the **Password** field.
3. Enter a profile name to be used for the system in the **Profile Name (optional)** field. This is the name under which the system will appear when viewed in Red Hat Network.
4. Select the method by which to register the hypervisor:
  - **The Red Hat Network**  
Select the **RHN** option and press the space bar to register the hypervisor directly with the Red Hat Network. You do not need to enter values in the **URL** and **CA URL** fields.

**Example 7.13. Red Hat Network Configuration**

RHN    Satellite    SAM

URL: \_\_\_\_\_

CA URL: \_\_\_\_\_

- **Satellite**

- a. Select the **Satellite** option and press the space bar to register the hypervisor with a Satellite server.
- b. Enter the URL of the Satellite server in the **URL** field.
- c. Enter the URL of the certificate authority for the Satellite server in the **CA URL** field.

**Example 7.14. Satellite Configuration**

RHN    Satellite    SAM

RHN URL: https://your-satellite.example.com\_\_\_\_\_

CA URL: https://your-satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT

- **Subscription Asset Manager**

- a. Select the **Subscription Asset Manager** option and press **Space** to register the hypervisor via Subscription Asset Manager.
- b. Enter the URL of the Subscription Asset Manager server in the **URL** field.
- c. Enter the URL of the certificate authority for the Subscription Asset Manager server in the **CA URL** field.

**Example 7.15. Subscription Asset Manager Configuration**

RHN    Satellite    SAM

URL: https://subscription-asset-manager.example.com\_\_\_\_\_

CA : https://subscription-asset-manager.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT

5. If you are using a proxy server, you must also specify the details of that server:
  - a. Enter the IP address or fully qualified domain name of the proxy server in the **Server** field.
  - b. Enter the port by which to attempt a connection to the proxy server in the **Port** field.
  - c. Enter the user name by which to attempt a connection to the proxy server in the **Username** field.

d. Enter the password by which to authenticate the user name specified above in the **Password** field.

6. Select **<Save>** and press **Enter**.

### Result

You have registered the hypervisor directly with the Red Hat Network, via a Satellite server or via SubSubscription Asset Manager.

[Report a bug](#)

## 7.6. ADDING HYPERVISORS TO RED HAT ENTERPRISE VIRTUALIZATION MANAGER

### 7.6.1. Using the Hypervisor

If the Hypervisor was configured with the address of the Red Hat Enterprise Virtualization Manager, the Hypervisor reboots and is automatically registered with the Manager. The Red Hat Enterprise Virtualization Manager interface displays the Hypervisor under the **Hosts** tab. To prepare the Hypervisor for use, it must be approved using Red Hat Enterprise Virtualization Manager.

If the Hypervisor was configured without the address of the Red Hat Enterprise Virtualization Manager, it must be added manually. To add the Hypervisor manually, you must have both the IP address of the machine upon which it was installed and the password that was set on the **oVirt Engine** screen during configuration.

[Report a bug](#)

### 7.6.2. Approving a Hypervisor

#### Summary

It is not possible to run virtual machines on a Hypervisor until the addition of it to the environment has been approved in Red Hat Enterprise Virtualization Manager.

#### Procedure 7.24. Approving a Hypervisor

1. Log in to the Red Hat Enterprise Virtualization Manager Administration Portal.
2. From the **Hosts** tab, click on the host to be approved. The host should currently be listed with the status of **Pending Approval**.
3. Click the **Approve** button. The **Edit and Approve Hosts** dialog displays. You can use the dialog to set a name for the host, fetch its SSH fingerprint before approving it, and configure power management, where the host has a supported power management card. For information on power management configuration, refer to [Section 9.8.2, "Host Power Management Settings Explained"](#).
4. Click **OK**. If you have not configured power management you will be prompted to confirm that you wish to proceed without doing so, click **OK**.

### Result

The status in the **Hosts** tab changes to **Installing**, after a brief delay the host status changes to **Up**.

[Report a bug](#)

## 7.7. MODIFYING THE RED HAT ENTERPRISE VIRTUALIZATION HYPERVISOR ISO

### 7.7.1. Introduction to Modifying the Red Hat Enterprise Virtualization Hypervisor ISO

While the Red Hat Enterprise Virtualization Hypervisor is designed as a closed, minimal operating system, you can use the **edit-node** tool to make specific changes to the Red Hat Enterprise Virtualization Hypervisor ISO file to address specific requirements. The tool extracts the file system from a **livecd**-based ISO file and modifies aspects of the image, such as user passwords, SSH keys, and the packages included.



#### IMPORTANT

Any modifications must be repeated each time prior to upgrading a hypervisor to a new version of the Red Hat Enterprise Virtualization Hypervisor ISO file.



#### WARNING

In the event of an issue with a Red Hat Enterprise Virtualization Hypervisor that has been modified using the **edit-node** tool, you may be required to reproduce the issue in an unmodified version of the Red Hat Enterprise Virtualization Hypervisor as part of the troubleshooting process.

[Report a bug](#)

### 7.7.2. Installing the edit-node Tool

#### Summary

The **edit-node** tool is included in the `ovirt-node-tools` package provided by the Red Hat Enterprise Virtualization Hypervisor channel.

#### Procedure 7.25. Installing the edit-node Tool

1. Log in to the system on which to modify the Red Hat Enterprise Virtualization Hypervisor ISO file.
2. Enable the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86\_64)** repository:

- With RHN Classic:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev
```

- With Subscription Manager, attach a **Red Hat Enterprise Virtualization** entitlement and run the following command:

```
# subscription-manager repos --enable=rhel-6-server-rhev-rpms
```

3. Install the `ovirt-node-tools` package:

```
# yum install ovirt-node-tools
```

## Result

You have installed the **edit-node** tool required for modifying the Red Hat Enterprise Virtualization Hypervisor ISO file.

[Report a bug](#)

### 7.7.3. Syntax of the edit-node Tool

The basic options for the **edit-node** tool are as follows:

#### Options for the edit-node Tool

##### **--name=*image\_name***

Specifies the name of the modified image.

##### **--output=*directory***

Specifies the directory to which the edited ISO is saved.

##### **--kickstart=*kickstart\_file***

Specifies the path or URL to and name of a kickstart configuration file.

##### **--script=*script***

Specifies the path to and name of a script to run in the image.

##### **--shell**

Opens an interactive shell with which to edit the image.

##### **--passwd=*user,encrypted\_password***

Defines a password for the specified *user*. This option accepts MD5-encrypted password values. The **--password** parameter can be specified multiple times to modify multiple users. If no user is specified, the default user is **admin**.

##### **--sshkey=*user,public\_key\_file***

Specifies the public key for the specified user. This option can be specified multiple times to specify keys for multiple users. If no user is specified, the default user is **admin**.

##### **--uidmod=*user,uid***

Specifies the user ID for the specified user. This option can be specified multiple times to specify IDs for multiple users.

##### **--gidmod=*group,gid***

Specifies the group ID for the specified group. This option can be specified multiple times to specify IDs for multiple groups.

**--tmpdir=*temporary\_directory***

Specifies the temporary directory on the local file system to use. By default, this value is set to **/var/tmp**

**--releasefile=*release\_file***

Specifies the path to and name of a release file to use for branding.

**--builder=*builder***

Specifies the builder of a remix.

**--install-plugin=*plugin***

Specifies a list of plug-ins to install in the image. You can specify multiple plug-ins by separating the plug-in names using a comma.

**--install=*package***

Specifies a list of packages to install in the image. You can specify multiple packages by separating the package names using a comma.

**--install-kmod=*package\_name***

Installs the specified driver update package from a **yum** repository or specified **.rpm** file. Specified **.rpm** files are valid only if in whitelisted locations (kmod-specific areas).

**--repo=*repository***

Specifies the yum repository to be used in conjunction with the **--install-\*** options. The value specified can be a local directory, a yum repository file (**.repo**), or a driver disk **.iso** file.

**--nogpgcheck**

Skips GPG key verification during the **yum install** stage. This option allows you to install unsigned packages.

## Manifest Options for theedit-node Tool

**--list-plugins**

Prints a list of plug-ins added to the image.

**--print-version**

Prints current version information from **/etc/system-release**.

**--print-manifests**

Prints a list of manifest files in the ISO file.

**--print-manifest=*manifest***

Prints the specified manifest file.

**--get-manifests=*manifest***

Creates a **.tar** file of manifest files in the ISO file.

**--print-file-manifest**

Prints the contents of **rootfs** on the ISO file.

**--print-rpm-manifest**

Prints a list of installed packages in **rootfs** on the ISO file.

**Debugging Options for the `edit-node` Tool****--debug**

Prints debugging information when the **edit-node** command is run.

**--verbose**

Prints verbose information regarding the progress of the **edit-node** command.

**--logfile=*logfile***

Specifies the path to and name of a file in which to print debugging information.

[Report a bug](#)

**7.7.4. Adding and Updating Packages**

You can use the **edit-node** tool to add new packages to or update existing packages in the Red Hat Enterprise Virtualization Hypervisor ISO file. To add or update a single package, you must either set up a local directory to act as a repository for the required package and its dependencies or point the **edit-node** tool to the location of a repository definition file that defines one or more repositories that provide the package and its dependencies. To add or update multiple packages, you must point the **edit-node** tool to the location of a repository definition file that defines one or more repositories that provide the packages and their dependencies.

**NOTE**

If you include a definition for a local repository in a repository definition file, the directory that acts as the source for that repository must be exposed via a web server or an FTP server. For example, it must be possible to access the repository via a link such as **http://localhost/myrepo/** or **ftp://localhost/myrepo/**.

**IMPORTANT**

The **edit-node** tool cannot download packages from repositories that use SSL. Instead, you must manually download each package and its dependencies and create a local repository that contains those packages.

[Report a bug](#)

**7.7.4.1. Creating a Local Repository**

## Summary

To add packages to the Red Hat Enterprise Virtualization Hypervisor ISO file, you must set up a directory to act as a repository for installing those packages using the **createrepo** tool provided by the base Red Hat Enterprise Linux Workstation and Red Hat Enterprise Linux Server channels.

### Procedure 7.26. Creating a Local Repository

1. Install the **createrepo** package and dependencies on the system on which to modify the Red Hat Enterprise Virtualization Hypervisor ISO file:

```
# yum install createrepo
```

2. Create a directory to serve as the repository.
3. Copy all required packages and their dependencies into the newly created directory.
4. Set up the metadata files for that directory to act as a repository:

```
# createrepo [directory_name]
```

## Result

You have created a local repository for installing the required packages and their dependencies in the Red Hat Enterprise Virtualization Hypervisor ISO file.

[Report a bug](#)

### 7.7.4.2. Example: Adding Packages to the Red Hat Enterprise Virtualization Hypervisor ISO File

You can use the **edit-node** tool to add packages to the Red Hat Enterprise Virtualization Hypervisor ISO file. This action creates a copy of the ISO file in the directory from which the **edit-node** tool was run that includes the name of the newly added packages in its name.

The following example adds a single package to the Red Hat Enterprise Virtualization Hypervisor ISO file, using a directory configured to act as a local repository as the source from which to install the package:

#### Example 7.16. Adding a Single Package to the Red Hat Enterprise Virtualization Hypervisor ISO File

```
# edit-node --nogpgcheck --install package1 --repo ./local_repo /usr/share/rhev-hypervisor/rhev-latest-6.iso
```

You can add multiple packages by enclosing a comma-separated list of package names in double quotation marks. The following example adds two packages to the Red Hat Enterprise Virtualization Hypervisor ISO file, using a directory configured to act as a local repository as the source from which to install the packages:

#### Example 7.17. Adding Multiple Packages to the Red Hat Enterprise Virtualization Hypervisor ISO File



```
# edit-node --nogpgcheck --install "package1,package2" --repo ./local_repo /usr/share/rhev-
hypervisor/rhev-latest-6.iso
```

[Report a bug](#)

### 7.7.4.3. Example: Updating Packages in the Red Hat Enterprise Virtualization Hypervisor ISO File

You can use the **edit-node** tool to update existing packages in the Red Hat Enterprise Virtualization Hypervisor ISO file. This action creates a copy of the ISO file in the directory from which the **edit-node** tool was run that includes the names of the updated packages in its name.

The following example updates the `vdsm` package in the Red Hat Enterprise Virtualization Hypervisor ISO file, using a repository file containing the details of the Red Hat Enterprise Virtualization Hypervisor repository:

#### Example 7.18. Updating a Single Package in the Red Hat Enterprise Virtualization Hypervisor ISO File

```
# edit-node --nogpgcheck --install vdsm --repo /etc/yum.repos.d/rhev.repo /usr/share/rhev-
hypervisor/rhev-latest-6.iso
```

You can update multiple packages by enclosing a comma-separated list of package names in double quotation marks. The following example updates the `vdsm` and `libvirt` packages in the Red Hat Enterprise Virtualization Hypervisor ISO file, using a repository file containing the details of the Red Hat Enterprise Virtualization Hypervisor repository:

#### Example 7.19. Updating Multiple Packages in the Red Hat Enterprise Virtualization Hypervisor ISO File

```
# edit-node --nogpgcheck --install "vdsm,libvirt" --repo /etc/yum.repos.d/rhev.repo
/usr/share/rhev-hypervisor/rhev-latest-6.iso
```

[Report a bug](#)

## 7.7.5. Modifying the Default ID of Users and Groups

### 7.7.5.1. Example: Modifying the Default ID of a User

You can use the **edit-node** tool to modify the default ID of a user in the Red Hat Enterprise Virtualization Hypervisor ISO file.

The following example changes the default ID of the user **user1** to **60**:

#### Example 7.20. Modifying the Default ID of a Single User

```
# edit-node --uidmod=user1,60
```

You can modify the default ID of multiple users by specifying the **--uidmod** option multiple times in the same command. The following example changes the default ID of the user **user1** to **60** and the default ID of the user **user2** to **70**.

#### Example 7.21. Modifying the Default ID of Multiple Users

```
# edit-node --uidmod=user1,60 --uidmod=user2,70
```

[Report a bug](#)

### 7.7.5.2. Example: Modifying the Default ID of a Group

You can use the **edit-node** tool to modify the default ID of a group in the Red Hat Enterprise Virtualization Hypervisor ISO file.

The following example changes the default ID of the group **group1** to **60**:

#### Example 7.22. Modifying the Default ID of a Single Group

```
# edit-node --gidmod=group1,60
```

You can modify the default ID of multiple groups by specifying the **--gidmod** option multiple times in the same command. The following example changes the default ID of the group **group1** to **60** and the default ID of the group **group2** to **70**.

#### Example 7.23. Modifying the Default ID of Multiple Groups

```
# edit-node --gidmod=group1,60 --gidmod=group2,70
```

[Report a bug](#)

## CHAPTER 8. RED HAT ENTERPRISE LINUX HOSTS

### 8.1. RED HAT ENTERPRISE LINUX HOSTS

You can use a standard Red Hat Enterprise Linux 6 installation on capable hardware as a host. Red Hat Enterprise Virtualization supports hosts running Red Hat Enterprise Linux 6 Server AMD64/Intel 64 version.

Adding a host 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 hand-shake process as the host and management system establish a connection.

[Report a bug](#)

### 8.2. HOST COMPATIBILITY MATRIX

Red Hat Enterprise Linux Version	Red Hat Enterprise Virtualization 3.4 clusters with 3.0 compatibility level	Red Hat Enterprise Virtualization 3.4 clusters with 3.1 compatibility level	Red Hat Enterprise Virtualization 3.4 clusters with 3.2 compatibility level	Red Hat Enterprise Virtualization 3.4 clusters with 3.3 compatibility level	Red Hat Enterprise Virtualization 3.4 clusters
6.2	Supported	Unsupported	Unsupported	Unsupported	Unsupported
6.3	Supported	Supported	Unsupported	Unsupported	Unsupported
6.4	Supported	Supported	Supported	Unsupported	Unsupported
6.5	Supported	Supported	Supported	Supported	Supported

[Report a bug](#)

## PART IV. BASIC SETUP

## CHAPTER 9. CONFIGURING HOSTS

### 9.1. INSTALLING RED HAT ENTERPRISE LINUX


#### Summary

You must install Red Hat Enterprise Linux Server 6.5 or 6.6 on a system to use it as a virtualization host in a Red Hat Enterprise Virtualization 3.4 environment.

#### Procedure 9.1. Installing Red Hat Enterprise Linux

##### 1. Download and Install Red Hat Enterprise Linux

Download and Install Red Hat Enterprise Linux Server 6.5 or 6.6 on the target virtualization host, referring to the *Red Hat Enterprise Linux Installation Guide* for detailed instructions. Only the Base package group is required to use the virtualization host in a Red Hat Enterprise Virtualization environment, though the host must be registered and subscribed to a number of entitlements before it can be added to the Manager.



#### IMPORTANT

If you intend to use directory services for authentication on the Red Hat Enterprise Linux host then you must ensure that the authentication files required by the **useradd** command are locally accessible. The `vdsmd` package, which provides software that is required for successful connection to Red Hat Enterprise Virtualization Manager, will not install correctly if these files are not locally accessible.

##### 2. Ensure Network Connectivity

Following successful installation of Red Hat Enterprise Linux Server 6.5 or 6.6, ensure that there is network connectivity between your new Red Hat Enterprise Linux host and the system on which your Red Hat Enterprise Virtualization Manager is installed.

- a. Attempt to ping the Manager:

```
# ping address of manager
```

- b. ■ If the Manager can successfully be contacted, this displays:

```
ping manager.example.com
PING manager.example.redhat.com (192.168.0.1) 56(84) bytes of data.
64 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.415 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.419 ms
64 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=1.41 ms
64 bytes from 192.168.0.1: icmp_seq=4 ttl=64 time=0.487 ms
64 bytes from 192.168.0.1: icmp_seq=5 ttl=64 time=0.409 ms
64 bytes from 192.168.0.1: icmp_seq=6 ttl=64 time=0.372 ms
64 bytes from 192.168.0.1: icmp_seq=7 ttl=64 time=0.464 ms

--- manager.example.redhat.com ping statistics ---
7 packets transmitted, 7 received, 0% packet loss, time 6267ms
```

- If the Manager cannot be contacted, this displays:

```
ping: unknown host manager.example.com
```

You must configure the network so that the host can contact the Manager. First, disable **NetworkManager**. Then configure the networking scripts so that the host will acquire an ip address on boot.

- i. Disable **NetworkManager**.

```
# service NetworkManager stop
# chkconfig NetworkManager disable
```

- ii. Edit `/etc/sysconfig/network-scripts/ifcfg-eth0`. Find this line:

```
ONBOOT=no
```

Change that line to this:

```
ONBOOT=yes
```

- iii. Reboot the host machine.
- iv. Ping the Manager again:

```
# ping address of manager
```

If the host still cannot contact the Manager, it is possible that your host machine is not acquiring an IP address from DHCP. Confirm that DHCP is properly configured and that your host machine is properly configured to acquire an IP address from DHCP.

If the Manager can successfully be contacted, this displays:

```
ping manager.example.com
PING manager.example.redhat.com (192.168.0.1) 56(84) bytes of data.
64 bytes from 192.168.0.1: icmp_seq=1 ttl=64 time=0.415 ms
64 bytes from 192.168.0.1: icmp_seq=2 ttl=64 time=0.419 ms
64 bytes from 192.168.0.1: icmp_seq=3 ttl=64 time=1.41 ms
64 bytes from 192.168.0.1: icmp_seq=4 ttl=64 time=0.487 ms
64 bytes from 192.168.0.1: icmp_seq=5 ttl=64 time=0.409 ms
64 bytes from 192.168.0.1: icmp_seq=6 ttl=64 time=0.372 ms
64 bytes from 192.168.0.1: icmp_seq=7 ttl=64 time=0.464 ms

--- manager.example.com ping statistics ---
7 packets transmitted, 7 received, 0% packet loss, time 6267ms
```

## Result

You have installed Red Hat Enterprise Linux Server 6.5 or 6.6. You must complete additional configuration tasks before adding the virtualization host to your Red Hat Enterprise Virtualization environment.

[Report a bug](#)

## 9.2. SUBSCRIBING TO REQUIRED CHANNELS USING SUBSCRIPTION MANAGER

### Summary

To be used as a virtualization host, a Red Hat Enterprise Linux host must be registered and subscribed to a number of entitlements using either Subscription Manager or RHN Classic. You must follow the steps in this procedure to register and subscribe using Subscription Manager. Completion of this procedure will mean that you have:

- Registered the virtualization host to Red Hat Network using Subscription Manager.
- Attached the **Red Hat Enterprise Linux Server** entitlement to the virtualization host.
- Attached the **Red Hat Enterprise Virtualization** entitlement to the virtualization host.

Do *not* follow the steps in this procedure to register and subscribe using RHN Classic.

### Procedure 9.2. Subscribing to Required Channels using Subscription Manager

#### 1. Register

Run the **subscription-manager** command with the **register** parameter to register the system with Red Hat Network. To complete registration successfully, you must supply your Red Hat Network **Username** and **Password** when prompted.

```
# subscription-manager register
```

#### 2. Identify Available Entitlement Pools

To attach the correct entitlements to the system, you must first locate the identifiers for the required entitlement pools. Use the **list** action of the **subscription-manager** to find these.

To identify available subscription pools for **Red Hat Enterprise Linux Server**, use the command:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Linux Server"
```

To identify available subscription pools for **Red Hat Enterprise Virtualization**, use the command:

```
# subscription-manager list --available | grep -A8 "Red Hat Enterprise Virtualization"
```

#### 3. Attach Entitlements to the System

Using the pool identifiers you located in the previous step, attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the system. Use the **attach** parameter of the **subscription-manager** command, replacing **[POOLID]** with each of the pool identifiers:

```
# subscription-manager attach --pool=[POOLID]
```

#### 4. Enable the Red Hat Enterprise Virtualization Management Agents Repository

Run the following command to enable the Red Hat Enterprise Virtualization Management Agents (RPMs) repository:

```
# subscription-manager repos --enable=rhel-6-server-rhev-mgmt-agent-rpms
```

■

## Result

You have registered the virtualization host to Red Hat Network and attached the required entitlements using Subscription Manager.

[Report a bug](#)

## 9.3. SUBSCRIBING TO REQUIRED CHANNELS USING RHN CLASSIC

### Summary

To be used as a virtualization host, a Red Hat Enterprise Linux host must be registered and subscribed to a number of entitlements using either Subscription Manager or RHN Classic. You must follow the steps in this procedure to register and subscribe using RHN Classic. Completion of this procedure will mean that you have:

- Registered the virtualization host to Red Hat Network using RHN Classic.
- Subscribed the virtualization host to the **Red Hat Enterprise Linux Server (v. 6 for 64-bit AMD64 / Intel64)** channel.
- Subscribed the virtualization host to the **Red Hat Enterprise Virt Management Agent (v 6 x86\_64)** channel.

Do *not* follow the steps in this procedure to register and subscribe using Subscription Manager.

### Procedure 9.3. Subscribing to Required Channels using RHN Classic

#### 1. Register

If the machine is not already registered with Red Hat Network, run the **rhn\_register** command as **root** to register it. To complete registration successfully, you must supply your Red Hat Network **Username** and **Password**. Follow the prompts displayed by **rhn\_register** to complete registration of the system.

```
# rhn_register
```

#### 2. Subscribe to channels

You must subscribe the system to the required channels using either the web interface to Red Hat Network or the command line **rhn-channel** command.

##### o Using the Web Interface to Red Hat Network

To add a channel subscription to a system from the web interface:

- a. Log on to Red Hat Network (<http://rhn.redhat.com>).
- b. Move the mouse cursor over the **Subscriptions** link at the top of the screen, and then click the **Registered Systems** link in the menu that appears.
- c. Select the system to which you are adding channels from the list presented on the screen, by clicking the name of the system.
- d. Click **Alter Channel Subscriptions** in the **Subscribed Channels** section of the screen.
- e. Select the channels to be added from the list presented on the screen. To use the virtualization host in a Red Hat Enterprise Virtualization environment you must select:



- **Red Hat Enterprise Linux Server (v. 6 for 64-bit x86\_64)**; and
  - **Red Hat Enterprise Virt Management Agent (v 6 x86\_64)**.
- f. Click the **Change Subscription** button to finalize the change.
- **Using the `rhn-channel` command**  
Run the **`rhn-channel`** command to subscribe the virtualization host to each of the required channels. Use the commands:

```
# rhn-channel --add --channel=rhel-x86_64-server-6
# rhn-channel --add --channel=rhel-x86_64-rhev-mgmt-agent-6
```



### IMPORTANT

If you are not the administrator for the machine as defined in Red Hat Network, or the machine is not registered to Red Hat Network, then use of the **`rhn-channel`** command will result in an error:

```
Error communicating with server. The message was:Error Class Code: 37
Error Class Info: You are not allowed to perform administrative tasks on
this system.
```

Explanation:

An error has occurred while processing your request. If this problem persists please enter a bug report at [bugzilla.redhat.com](http://bugzilla.redhat.com).

If you choose to submit the bug report, please be sure to include details of what you were trying to do when this error occurred and details on how to reproduce this problem.

If you encounter this error when using **`rhn-channel`**, you must instead use the web interface to add the channel to the system.

### Result

You have registered the virtualization host to Red Hat Network and subscribed to the required entitlements using RHN Classic.

[Report a bug](#)

## 9.4. CONFIGURING VIRTUALIZATION HOST FIREWALL

### Summary

Red Hat Enterprise Virtualization requires that a number of network ports be open to support virtual machines and remote management of the virtualization host from the Red Hat Enterprise Virtualization Manager. You must follow this procedure to open the required network ports before attempting to add the virtualization host to the Manager.

### Procedure 9.4. Configuring Virtualization Host Firewall

The following steps configure the default firewall in Red Hat Enterprise Linux, **`iptables`**, to allow traffic on the required network ports. This procedure replaces the host's existing firewall configuration with one that contains only the ports required by Red Hat Enterprise Virtualization. If you have existing firewall

rules with which this configuration must be merged, then you must do so by manually editing the rules defined in the **iptables** configuration file, **/etc/sysconfig/iptables**.

All commands in this procedure must be run as the **root** user.

### 1. Remove existing firewall rules from configuration

Remove any existing firewall rules using the **--flush** parameter to the **iptables** command.

```
# iptables --flush
```

### 2. Add new firewall rules to configuration

Add the new firewall rules, required by Red Hat Enterprise Virtualization, using the **--append** parameter to the **iptables** command. The prompt character (**#**) has been intentionally omitted from this list of commands to allow easy copying of the content to a script file or command prompt.

```
iptables --append INPUT -m state --state ESTABLISHED,RELATED -j ACCEPT
iptables --append INPUT -p icmp -j ACCEPT
iptables --append INPUT -i lo -j ACCEPT
iptables --append INPUT -p tcp --dport 22 -j ACCEPT
iptables --append INPUT -p tcp --dport 16514 -j ACCEPT
iptables --append INPUT -p tcp --dport 54321 -j ACCEPT
iptables --append INPUT -p tcp -m multiport --dports 5634:6166 -j ACCEPT
iptables --append INPUT -p tcp -m multiport --dports 49152:49216 -j ACCEPT
iptables --append INPUT -j REJECT --reject-with icmp-host-prohibited
iptables --append FORWARD -m physdev ! --physdev-is-bridged -j REJECT \
--reject-with icmp-host-prohibited
```

#### NOTE

The provided **iptables** commands add firewall rules to accept network traffic on a number of ports. These include:

- port **22** for **SSH**,
- ports **5634** to **6166** for guest console connections,
- port **16514** for **libvirt** virtual machine migration traffic,
- ports **49152** to **49216** for **VDSM** virtual machine migration traffic, and
- port **54321** for the Red Hat Enterprise Virtualization Manager.

### 3. Save the updated firewall configuration

Save the updated firewall configuration script using the **save** to the **iptables** initialization script.

```
# service iptables save
```

### 4. Enable iptables service

Ensure that the **iptables** service is configured to start on boot and has been restarted, or started for the first time if it was not already running.

```
# chkconfig iptables on
# service iptables restart
```

### Result

You have configured the virtualization host's firewall to allow the network traffic required by Red Hat Enterprise Virtualization.

[Report a bug](#)

## 9.5. CONFIGURING VIRTUALIZATION HOST SUDO

### Summary

The Red Hat Enterprise Virtualization Manager uses **sudo** to perform operations as the **root** on the host. The default Red Hat Enterprise Linux configuration, stored in **/etc/sudoers**, contains values that allow this. If this file has been modified since Red Hat Enterprise Linux installation, then these values may have been removed. This procedure verifies that the required entry still exists in the configuration, and adds the required entry if it is not present.

### Procedure 9.5. Configuring Virtualization Host sudo

1. **Log in**

Log in to the virtualization host as the **root** user.

2. **Run visudo**

Run the **visudo** command to open the **/etc/sudoers** file.

```
# visudo
```

3. **Edit sudoers file**

Read the configuration file, and verify that it contains these lines:

```
# Allow root to run any commands anywhere
root ALL=(ALL) ALL
```

If the file does not contain these lines, add them and save the file using the VIM **:w** command.

4. **Exit editor**

Exit **visudo** using the VIM **:q** command.

### Result

You have configured **sudo** to allow use by the **root** user.

[Report a bug](#)

## 9.6. CONFIGURING VIRTUALIZATION HOST SSH

### Summary

The Red Hat Enterprise Virtualization Manager accesses virtualization hosts via SSH. To do this it logs in as the **root** user using an encrypted key for authentication. You must follow this procedure to ensure that SSH is configured to allow this.

**WARNING**

The first time the Red Hat Enterprise Virtualization Manager is connected to the host it will install an authentication key. In the process it will overwrite any existing keys contained in the **/root/.ssh/authorized\_keys** file.

**Procedure 9.6. Configuring virtualization host SSH**

All commands in this procedure must be run as the **root** user.

**1. Install the SSH server (openssh-server)**

Install the openssh-server package using **yum**.

```
# yum install openssh-server
```

**2. Edit SSH server configuration**

Open the SSH server configuration, **/etc/ssh/sshd\_config**, in a text editor. Search for the **PermitRootLogin**.

- If **PermitRootLogin** is set to **yes**, or is not set at all, no further action is required.
- If **PermitRootLogin** is set to **no**, then you must change it to **yes**.

Save any changes that you have made to the file, and exit the text editor.

**3. Enable the SSH server**

Configure the SSH server to start at system boot using the **chkconfig** command.

```
# chkconfig --level 345 sshd on
```

**4. Start the SSH server**

Start the SSH, or restart it if it is already running, using the **service** command.

```
# service sshd restart
```

**Result**

You have configured the virtualization host to allow **root** access over SSH.

[Report a bug](#)

**9.7. ADDING A RED HAT ENTERPRISE LINUX HOST****Summary**

A Red Hat Enterprise Linux host is based on a standard "basic" installation of Red Hat Enterprise Linux. The physical host must be set up before you can add it to the Red Hat Enterprise Virtualization environment.


The Red Hat Enterprise Virtualization Manager logs into the host to perform virtualization capability checks, install packages, create a network bridge, and reboot the host. The process of adding a new host can take up to 10 minutes.

### Procedure 9.7. Adding a Red Hat Enterprise Linux Host

1. Click the **Hosts** resource tab to list the hosts in the results list.
2. Click **New** to open the **New Host** window.
3. Use the drop-down menus to select the **Data Center** and **Host Cluster** for the new host.
4. Enter the **Name**, **Address**, and **SSH Port** of the new host.
5. Select an authentication method to use with the host.
  - Enter the root user's password to use password authentication.
  - Copy the key displayed in the **SSH PublicKey** field to `/root/.ssh/authorized_keys` on the host to use public key authentication.
6. You have now completed the mandatory steps to add a Red Hat Enterprise Linux host. 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. You can configure the **Power Management** and **SPM** using the applicable tabs now; however, as these are not fundamental to adding a Red Hat Enterprise Linux host, they are not covered in this procedure.
8. Click **OK** to add the host and close the window.

### Result

The new host displays in the list of hosts with a status of **Installing**. When installation is complete, the status updates to **Reboot**. The host must be activated for the status to change to **Up**.



#### NOTE

You can view the progress of the installation in the details pane.

[Report a bug](#)

## 9.8. EXPLANATION OF SETTINGS AND CONTROLS IN THE NEW HOST AND EDIT HOST WINDOWS

### 9.8.1. Host General Settings Explained

These settings apply when editing the details of a host or adding new Red Hat Enterprise Linux hosts and Foreman host provider hosts.

The **General** settings table contains the information required on the **General** tab of the **New Host** or **Edit Host** window.

Table 9.1. General settings

Field Name	Description
<b>Data Center</b>	The data center to which the host belongs. Red Hat Enterprise Virtualization Hypervisor hosts cannot be added to Gluster-enabled clusters.
<b>Host Cluster</b>	The cluster to which the host belongs.
<b>Use External Providers</b>	<p>Select or clear this check box to view or hide options for adding hosts provided by external providers. Upon selection, a drop-down list of external providers that have been added to the Manager displays. The following options are also available:</p> <ul style="list-style-type: none"> <li>● <b>Provider search filter</b> - A text field that allows you to search for hosts provided by the selected external provider. This option is provider-specific; see provider documentation for details on forming search queries for specific providers. Leave this field blank to view all available hosts.</li> <li>● <b>External Hosts</b> - A drop-down list that is populated with the name of hosts provided by the selected external provider. The entries in this list are filtered in accordance with any search queries that have been input in the <b>Provider search filter</b>.</li> </ul>
<b>Name</b>	The name of the cluster. This text field has a 40-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Comment</b>	A field for adding plain text, human-readable comments regarding the host.
<b>Address</b>	The IP address, or resolvable hostname of the host.
<b>Password</b>	The password of the host's root user. This can only be given when you add the host; it cannot be edited afterwards.
<b>SSH PublicKey</b>	Copy the contents in the text box to the <b>/root/.known_hosts</b> file on the host to use the Manager's ssh key instead of using a password to authenticate with the host.
<b>Automatically configure host firewall</b>	When adding a new host, the Manager can open the required ports on the host's firewall. This is enabled by default. This is an <b>Advanced Parameter</b> .

Field Name	Description
<b>SSH Fingerprint</b>	You can <b>fetch</b> the host's SSH fingerprint, and compare it with the fingerprint you expect the host to return, ensuring that they match. This is an <b>Advanced Parameter</b> .

[Report a bug](#)

## 9.8.2. Host Power Management Settings Explained

The **Power Management** settings table contains the information required on the **Power Management** tab of the **New Host** or **Edit Host** windows.

**Table 9.2. Power Management Settings**

Field Name	Description
<b>Primary/ Secondary</b>	Prior to Red Hat Enterprise Virtualization 3.2, a host with power management configured only recognized one fencing agent. Fencing agents configured on version 3.1 and earlier, and single agents, are treated as primary agents. The secondary option is valid when a second agent is defined.
<b>Concurrent</b>	Valid when there are two fencing agents, for example for dual power hosts in which each power switch has two agents connected to the same power switch. <ul style="list-style-type: none"> <li>● If this check box is selected, both fencing agents are used concurrently when a host is fenced. This means that both fencing agents have to respond to the Stop command for the host to be stopped; if one agent responds to the Start command, the host will go up.</li> <li>● If this check box is not selected, the fencing agents are used sequentially. This means that to stop or start a host, the primary agent is used first, and if it fails, the secondary agent is used.</li> </ul>
<b>Address</b>	The address to access your host's power management device. Either a resolvable hostname or an IP address.
<b>User Name</b>	User account with which to access the power management device. You can set up a user on the device, or use the default user.
<b>Password</b>	Password for the user accessing the power management device.

Field Name	Description
<b>Type</b>	<p>The type of power management device in your host.</p> <p>Choose one of the following:</p> <ul style="list-style-type: none"> <li>● <b>apc</b> - APC MasterSwitch network power switch. Not for use with APC 5.x power switch devices.</li> <li>● <b>apc_snmp</b> - Use with APC 5.x power switch devices.</li> <li>● <b>bladecenter</b> - IBM Bladecentre Remote Supervisor Adapter.</li> <li>● <b>cisco_ucs</b> - Cisco Unified Computing System.</li> <li>● <b>drac5</b> - Dell Remote Access Controller for Dell computers.</li> <li>● <b>drac7</b> - Dell Remote Access Controller for Dell computers.</li> <li>● <b>eps</b> - ePowerSwitch 8M+ network power switch.</li> <li>● <b>hpblade</b> - HP BladeSystem.</li> <li>● <b>ilo, ilo2, ilo3, ilo4</b> - HP Integrated Lights-Out.</li> <li>● <b>ipmilan</b> - Intelligent Platform Management Interface and Sun Integrated Lights Out Management devices.</li> <li>● <b>rsa</b> - IBM Remote Supervisor Adaptor.</li> <li>● <b>rsb</b> - Fujitsu-Siemens RSB management interface.</li> <li>● <b>wti</b> - WTI Network PowerSwitch.</li> </ul>
<b>Port</b>	The port number used by the power management device to communicate with the host.
<b>Options</b>	Power management device specific options. Enter these as 'key=value' or 'key'. See the documentation of your host's power management device for the options available.
<b>Secure</b>	Tick this check box to allow the power management device to connect securely to the host. This can be done via ssh, ssl, or other authentication protocols depending on and supported by the power management agent.



Field Name	Description
<b>Source</b>	Specifies whether the host will search within its <b>cluster</b> or <b>data center</b> for a fencing proxy. Use the <b>Up</b> and <b>Down</b> buttons to change the sequence in which the resources are used.
<b>Disable policy control of power management</b>	Power management is controlled by the <b>Cluster Policy</b> of the host's <b>cluster</b> . If power management is enabled and the defined low utilization value is reached, the Manager will power down the host machine, and restart it again when load balancing requires or there are not enough free hosts in the cluster. Tick this check box to disable policy control.

[Report a bug](#)

### 9.8.3. SPM Priority Settings Explained

The **SPM** settings table details the information required on the **SPM** tab of the **New Host** or **Edit Host** window.

Table 9.3. SPM settings

Field Name	Description
<b>SPM Priority</b>	Defines the likelihood that the host will be given the role of Storage Pool Manager (SPM). The options are <b>Low</b> , <b>Normal</b> , and <b>High</b> priority. Low priority means that there is a reduced likelihood of the host being assigned the role of SPM, and High priority means there is an increased likelihood. The default setting is Normal.

[Report a bug](#)

### 9.8.4. Host Console Settings Explained

The **Console** settings table details the information required on the **Console** tab of the **New Host** or **Edit Host** window.

Table 9.4. Console settings

Field Name	Description
------------	-------------

Field Name	Description
<b>Override display address</b>	Select this check box to override the display addresses of the host. This feature is useful in a case where the hosts are defined by internal IP and are behind a NAT firewall. When a user connects to a virtual machine from outside of the internal network, instead of returning the private address of the host on which the virtual machine is running, the machine returns a public IP or FQDN (which is resolved in the external network to the public IP).
<b>Display address</b>	The display address specified here will be used for all virtual machines running on this host. The address must be in the format of a fully qualified domain name or IP.

[Report a bug](#)

## CHAPTER 10. CONFIGURING DATA CENTERS

### 10.1. WORKFLOW PROGRESS - PLANNING YOUR DATA CENTER

		<b>Check System Requirements</b>
		<b>Install Red Hat Enterprise Virtualization Manager</b>
		<b>Data Collection Setup and Reports Installation*</b>
		<b>Install Virtualization Hosts</b>
		<b>Plan Your Data Center</b>
		<b>Network Setup*</b>
		<b>Storage Setup</b>

\* Optional task

#115226

[Report a bug](#)

### 10.2. PLANNING YOUR DATA CENTER

Successful planning is essential for a highly available, scalable Red Hat Enterprise Virtualization environment.

Although it is assumed that your solution architect has defined the environment before installation, the following considerations must be made when designing the system.

#### CPU

Virtual Machines must be distributed across hosts so that enough capacity is available to handle higher than average loads during peak processing. Average target utilization will be 50% of available CPU.

#### Memory

The Red Hat Enterprise Virtualization page sharing process overcommits up to 150% of physical memory for virtual machines. Therefore, allow for an approximately 30% overcommit.

## Networking

When designing the network, it is important to ensure that the volume of traffic produced by storage, remote connections and virtual machines is taken into account. As a general rule, allow approximately 50 MBps per virtual machine.

It is best practice to separate disk I/O traffic from end-user traffic, as this reduces the load on the Ethernet connection and reduces security vulnerabilities by isolating data from the visual stream. For Ethernet networks, it is suggested that bonds (802.3ad) are utilized to aggregate server traffic types.



### NOTE

It is possible to connect both the storage and Hypervisors via a single high performance switch. For this configuration to be effective, the switch must be able to provide 30 GBps on the backplane.

## High Availability

The system requires at least two hosts to achieve high availability. This redundancy is useful when performing maintenance or repairs.

[Report a bug](#)

## 10.3. DATA CENTERS IN RED HAT ENTERPRISE VIRTUALIZATION

The data center is the highest level container for all physical and logical resources within a managed virtual environment. The data center is a collection of clusters of hosts. It owns the logical network (that is, the defined subnets for management, guest network traffic, and storage network traffic) and the storage pool.

Red Hat Enterprise Virtualization contains a **Default** data center at installation. You can create new data centers that will also be managed from the single Administration Portal. For example, you may choose to have different data centers for different physical locations, business units, or for reasons of security. It is recommended that you do not remove the **Default** data center; instead, set up new appropriately named data centers.

The system administrator, as the superuser, can manage all aspects of the platform, that is, data centers, storage domains, users, roles, and permissions, by default; however, more specific administrative roles and permissions can be assigned to other users. For example, the enterprise may need a Data Center administrator for a specific data center, or a particular cluster may need an administrator. All system administration roles for physical resources have a hierarchical permission system. For example, a data center administrator will automatically have permission to manage all the objects in that data center - including storage domains, clusters, and hosts.

[Report a bug](#)

## 10.4. CREATING A NEW DATA CENTER

### Summary

This procedure creates a data center in your virtualization environment. The data center requires a functioning cluster, host, and storage domain to operate.

**NOTE**

The storage **Type** can be edited until the first storage domain is added to the data center. Once a storage domain has been added, the storage **Type** cannot be changed.

If you set the **Compatibility Version** as **3.1**, it cannot be changed to **3.0** at a later time; version regression is not allowed.

**Procedure 10.1. Creating a New Data Center**

1. Select the **Data Centers** resource tab to list all data centers in the results list.
2. Click **New** to open the **New Data Center** window.
3. Enter the **Name** and **Description** of the data center.
4. Select the storage **Type**, **Compatibility Version**, and **Quota Mode** of the data center from the drop-down menus.
5. Click **OK** to create the data center and open the **New Data Center - Guide Me** window.
6. The **Guide Me** window lists the entities that need to be configured for the data center. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the data center and clicking the **Guide Me** button.

**Result**

The new data center is added to the virtualization environment. It will remain **Uninitialized** until a cluster, host, and storage domain are configured for it; use **Guide Me** to configure these entities.

[Report a bug](#)

**10.5. CHANGING THE DATA CENTER COMPATIBILITY VERSION****Summary**

Red Hat Enterprise Virtualization data centers have a compatibility version. The compatibility version indicates the version of Red Hat Enterprise Virtualization that the data center is intended to be compatible with. All clusters in the data center must support the desired compatibility level.

**NOTE**

To change the data center compatibility version, you must have first updated all the clusters in your data center to a level that supports your desired compatibility level.

**Procedure 10.2. Changing the Data Center Compatibility Version**

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Data Centers** tab.
3. Select the data center to change from the list displayed. If the list of data centers is too long to filter visually then perform a search to locate the desired data center.
4. Click the **Edit** button.

5. Change the **Compatibility Version** to the desired value.
6. Click **OK**.

### Result

You have updated the compatibility version of the data center.



#### WARNING

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center. If you are upgrading the compatibility version from below 3.1 to a higher version, these storage domains will become unusable with versions older than 3.1.

[Report a bug](#)

## CHAPTER 11. CONFIGURING CLUSTERS

### 11.1. CLUSTERS IN RED HAT ENTERPRISE VIRTUALIZATION

A cluster is a collection of physical hosts that share similar characteristics and work together to provide computing resources in a highly available manner. In Red Hat Enterprise Virtualization the cluster must contain physical hosts that share the same storage domains and have the same type of CPU. Because virtual machines can be migrated across hosts in the same cluster, the cluster is the highest level at which power and load-sharing policies can be defined. The Red Hat Enterprise Virtualization platform contains a **Default** cluster in the **Default** data center at installation time.

Every cluster in the system must belong to a data center, and every host in the system must belong to a cluster. This enables the system to dynamically allocate a virtual machine to any host in the cluster, according to policies defined on the **Cluster** tab, thus maximizing memory and disk space, as well as virtual machine uptime.

At any given time, after a virtual machine runs on a specific host in the cluster, the virtual machine can be migrated to another host in the cluster using **Migrate**. This can be very useful when a host must be shut down for maintenance. The migration to another host in the cluster is transparent to the user, and the user continues working as usual. Note that a virtual machine cannot be migrated to a host outside the cluster.



#### NOTE

Red Hat Enterprise Virtualization 3.1 supports the use of clusters to manage Gluster storage bricks, in addition to virtualization hosts. To begin managing Gluster storage bricks, create a cluster with the **Enable Gluster Service** option selected. For further information on Gluster storage bricks, see the *Red Hat Storage Administration Guide*, available at [https://access.redhat.com/site/documentation/en-US/Red\\_Hat\\_Storage/](https://access.redhat.com/site/documentation/en-US/Red_Hat_Storage/).



#### NOTE

Red Hat Enterprise Virtualization supports **Memory Optimization** by enabling and tuning Kernel Same-page Merging (KSM) on the virtualization hosts in the cluster. For more information on KSM, see the *Red Hat Enterprise Linux 6 Virtualization Administration Guide*.

[Report a bug](#)

### 11.2. CREATING A NEW CLUSTER

#### Summary

A data center can contain multiple clusters, and a cluster can contain multiple hosts. All hosts in a cluster must be of the same CPU type (Intel or AMD). It is recommended that you create your hosts before you create your cluster to ensure CPU type optimization. However, you can configure the hosts at a later time using the **Guide Me** button.

#### Procedure 11.1. Creating a New Cluster

1. Select the **Clusters** resource tab.
2. Click **New** to open the **New Cluster** window.

3. Select the **Data Center** the cluster will belong to from the drop-down list.
4. Enter the **Name** and **Description** of the cluster.
5. Select the **CPU Name** and **Compatibility Version** from the drop-down lists. It is important to match the CPU processor family with the minimum CPU processor type of the hosts you intend to attach to the cluster, otherwise the host will be non-operational.
6. Select either the **Enable Virt Service** or **Enable Gluster Service** radio button to define whether the cluster will be populated with virtual machine hosts or with Gluster-enabled nodes. Note that you cannot add Red Hat Enterprise Virtualization Hypervisor hosts to a Gluster-enabled cluster.
7. Click the **Optimization** tab to select the memory page sharing threshold for the cluster, and optionally enable CPU thread handling and memory ballooning on the hosts in the cluster.
8. Click the **Cluster Policy** tab to optionally configure a cluster policy, scheduler optimization settings, enable trusted service for hosts in the cluster, and enable HA Reservation.
9. Click the **Resilience Policy** tab to select the virtual machine migration policy.
10. Click the **Console** tab to optionally override the global SPICE proxy, if any, and specify the address of a SPICE proxy for hosts in the cluster.
11. Click **OK** to create the cluster and open the **New Cluster - Guide Me** window.
12. The **Guide Me** window lists the entities that need to be configured for the cluster. Configure these entities or postpone configuration by clicking the **Configure Later** button; configuration can be resumed by selecting the cluster and clicking the **Guide Me** button.

## Result

The new cluster is added to the virtualization environment.

[Report a bug](#)

## 11.3. CHANGING THE CLUSTER COMPATIBILITY VERSION

### Summary

Red Hat Enterprise Virtualization clusters have a compatibility version. The cluster compatibility version indicates the features of Red Hat Enterprise Virtualization supported by all of the hosts in the cluster. The cluster compatibility is set according to the version of the least capable host operating system in the cluster.



### NOTE

To change the cluster compatibility version, you must have first updated all the hosts in your cluster to a level that supports your desired compatibility level.

### Procedure 11.2. Changing the Cluster Compatibility Version

1. Log in to the Administration Portal as the administrative user. By default this is the **admin** user.
2. Click the **Clusters** tab.



3. Select the cluster to change from the list displayed. If the list of clusters is too long to filter visually then perform a search to locate the desired cluster.
4. Click the **Edit** button.
5. Change the **Compatibility Version** to the desired value.
6. Click **OK** to open the **Change Cluster Compatibility Version** confirmation window.
7. Click **OK** to confirm.

### Result

You have updated the compatibility version of the cluster. Once you have updated the compatibility version of all clusters in a data center, then you are also able to change the compatibility version of the data center itself.



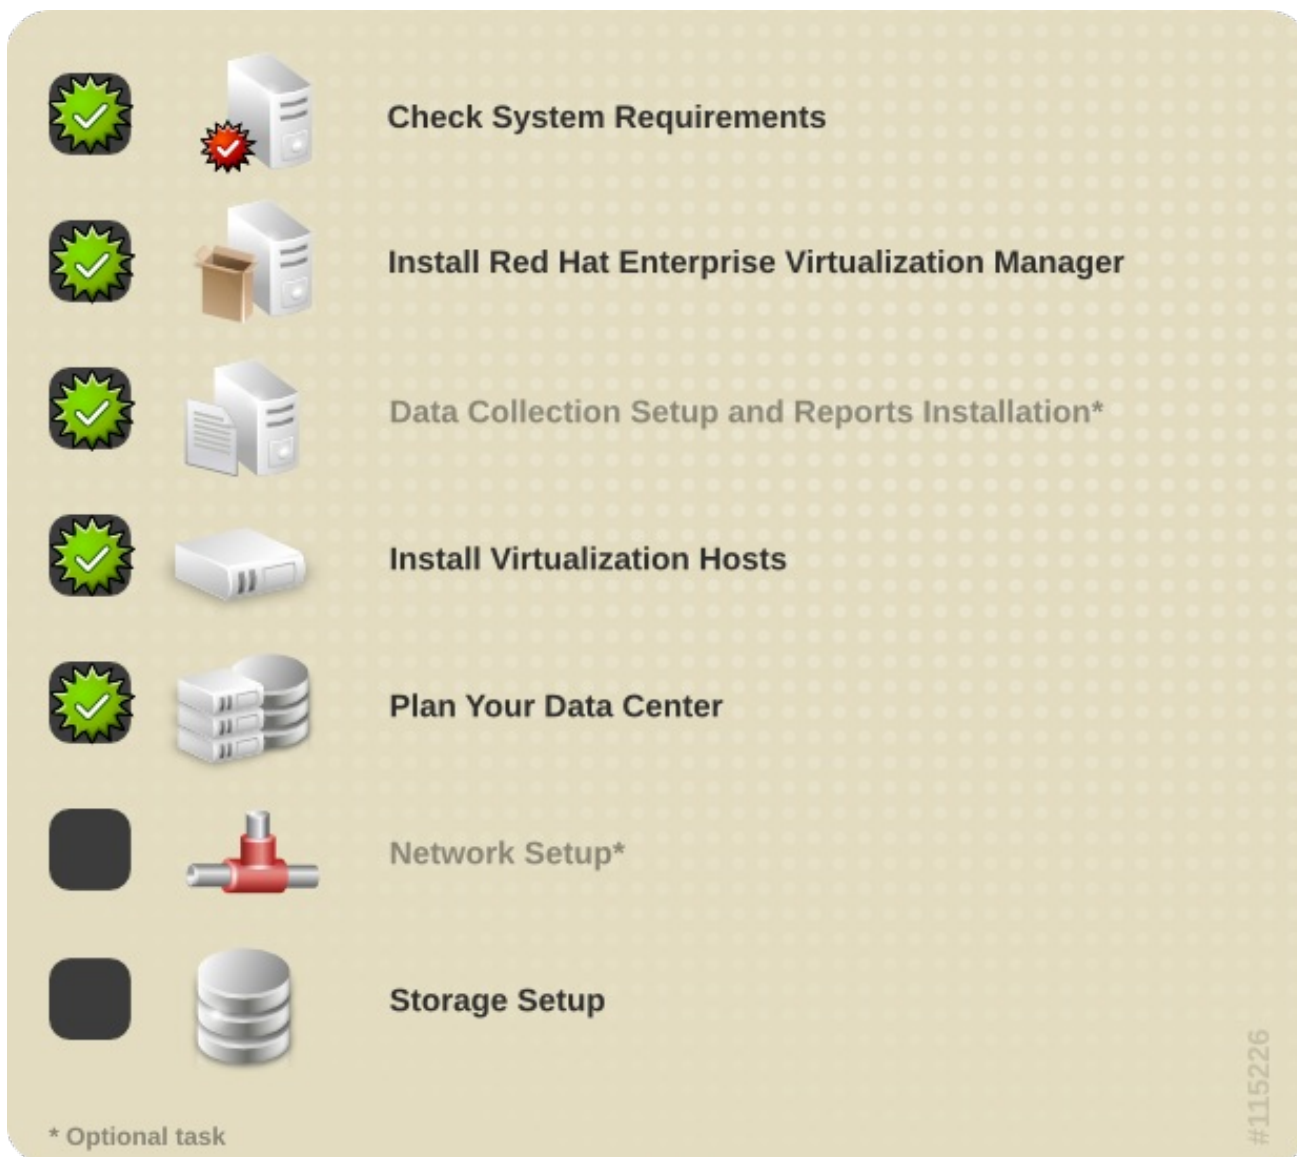
#### WARNING

Upgrading the compatibility will also upgrade all of the storage domains belonging to the data center. If you are upgrading the compatibility version from below 3.1 to a higher version, these storage domains will become unusable with versions older than 3.1.

[Report a bug](#)

## CHAPTER 12. CONFIGURING NETWORKING

### 12.1. WORKFLOW PROGRESS - NETWORK SETUP



[Report a bug](#)

### 12.2. NETWORKING IN RED HAT ENTERPRISE VIRTUALIZATION

Red Hat Enterprise Virtualization uses networking to support almost every aspect of operations. Storage, host management, user connections, and virtual machine connectivity, for example, all rely on a well planned and configured network to deliver optimal performance. Setting up networking is a vital prerequisite for a Red Hat Enterprise Virtualization environment because it is much simpler to plan for your projected networking requirements and implement your network accordingly than it is to discover your networking requirements through use and attempt to alter your network configuration retroactively.

It is however possible to deploy a Red Hat Enterprise Virtualization environment with no consideration given to networking at all. Simply ensuring that each physical machine in the environment has at least one *Network Interface Controller* (NIC) is enough to begin using Red Hat Enterprise Virtualization. While

it is true that this approach to networking will provide a functional environment, it will not provide an optimal environment. As network usage varies by task or action, grouping related tasks or functions into specialized networks can improve performance while simplifying the troubleshooting of network issues.

Red Hat Enterprise Virtualization separates network traffic by defining logical networks. Logical networks define the path that a selected network traffic type must take through the network. They are created to isolate network traffic by functionality or virtualize a physical topology.

The **rhevm** logical network is created by default and labeled as the **Management**. The **rhevm** logical network is intended for management traffic between the Red Hat Enterprise Virtualization Manager and virtualization hosts. You are able to define additional logical networks to segregate:

- Display related network traffic.
- General virtual machine network traffic.
- Storage related network traffic.

For optimal performance it is recommended that these traffic types be separated using logical networks. Logical networks may be supported using physical devices such as NICs or logical devices, such as network bonds. It is not necessary to have one device for each logical network as multiple logical networks are able to share a single device. This is accomplished using Virtual LAN (VLAN) tagging to isolate network traffic. To make use of this facility VLAN tagging must also be supported at the switch level.

The limits that apply to the number of logical networks that you may define in a Red Hat Enterprise Virtualization environment are:

- The number of logical networks attached to a host is limited to the number of available network devices combined with the maximum number of Virtual LANs (VLANs) which is 4096.
- The number of logical networks in a cluster is limited to the number of logical networks that can be attached to a host as networking must be the same for all hosts in a cluster.
- The number of logical networks in a data center is limited only by the number of clusters it contains in combination with the number of logical networks permitted per cluster.



#### NOTE

From Red Hat Enterprise Virtualization 3.3, network traffic for migrating virtual machines has been separated from network traffic for communication between the Manager and hosts. This prevents hosts from becoming non-responsive when importing or migrating virtual machines.




#### NOTE

A familiarity with the network concepts and their use is highly recommended when planning and setting up networking in a Red Hat Enterprise Virtualization environment. This document does not describe the concepts, protocols, requirements or general usage of networking. It is recommended that you read your network hardware vendor's guides for more information on managing networking.



## IMPORTANT

Additional care must be taken when modifying the properties of the **rhev** network. Incorrect changes to the properties of the **rhev** network may cause hosts to become temporarily unreachable.



## IMPORTANT

If you plan to use Red Hat Enterprise Virtualization nodes to provide any services, remember that the services will stop if the Red Hat Enterprise Virtualization environment stops operating.

This applies to all services, but you should be fully aware of the hazards of running the following on Red Hat Enterprise Virtualization:

- Directory Services
- DNS
- Storage

[Report a bug](#)

## 12.3. CREATING LOGICAL NETWORKS

### 12.3.1. Creating a New Logical Network in a Data Center or Cluster

#### Summary

Create a logical network and define its use in a data center, or in clusters in a data center.


#### Procedure 12.1. Creating a New Logical Network in a Data Center or Cluster

1. Use the **Data Centers** or **Clusters** resource tabs, tree mode, or the search function to find and select a data center or cluster in the results list.
2. Click the **Logical Networks** tab of the details pane to list the existing logical networks.
3. From the **Data Centers** details pane, click **New** to open the **New Logical Network** window.  
From the **Clusters** details pane, click **Add Network** to open the **New Logical Network** window.
4. Enter a **Name**, **Description** and **Comment** for the logical network.
5. In the **Export** section, select the **Create on external provider** check box to create the logical network on an external provider. Select the external provider from the **External Provider** drop-down menu.
6. In the **Network Parameters** section, select the **Enable VLAN tagging**, **VM network** and **Override MTU** to enable these options.
7. Enter a new label or select an existing label for the logical network in the **Network Label** text field.

8. From the **Cluster** tab, select the clusters to which the network will be assigned. You can also specify whether the logical network will be a required network.
9. From the **Subnet** tab, enter a **Name, CIDR** and select an **IP Version** for the subnet that the logical network will provide.
10. From the **Profiles** tab, add vNIC profiles to the logical network as required.
11. Click **OK**.

## Result

You have defined a logical network as a resource required by a cluster or clusters in the data center. If you entered a label for the logical network, it will be automatically added to all host network interfaces with that label.



## NOTE

When creating a new logical network or making changes to an existing logical network that is used as a display network, any running virtual machines that use that network must be rebooted before the network becomes available or the changes are applied.


[Report a bug](#)

## 12.4. EDITING LOGICAL NETWORKS

### 12.4.1. Editing Host Network Interfaces and Assigning Logical Networks to Hosts

#### Summary

You can change the settings of physical host network interfaces, move the management network from one physical host network interface to another, and assign logical networks to physical host network interfaces.



## IMPORTANT

You cannot assign logical networks offered by external providers to physical host network interfaces; such networks are dynamically assigned to hosts as they are required by virtual machines.

#### Procedure 12.2. Editing Host Network Interfaces and Assigning Logical Networks to Hosts

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results.
2. Click the **Network Interfaces** tab in the details pane.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.

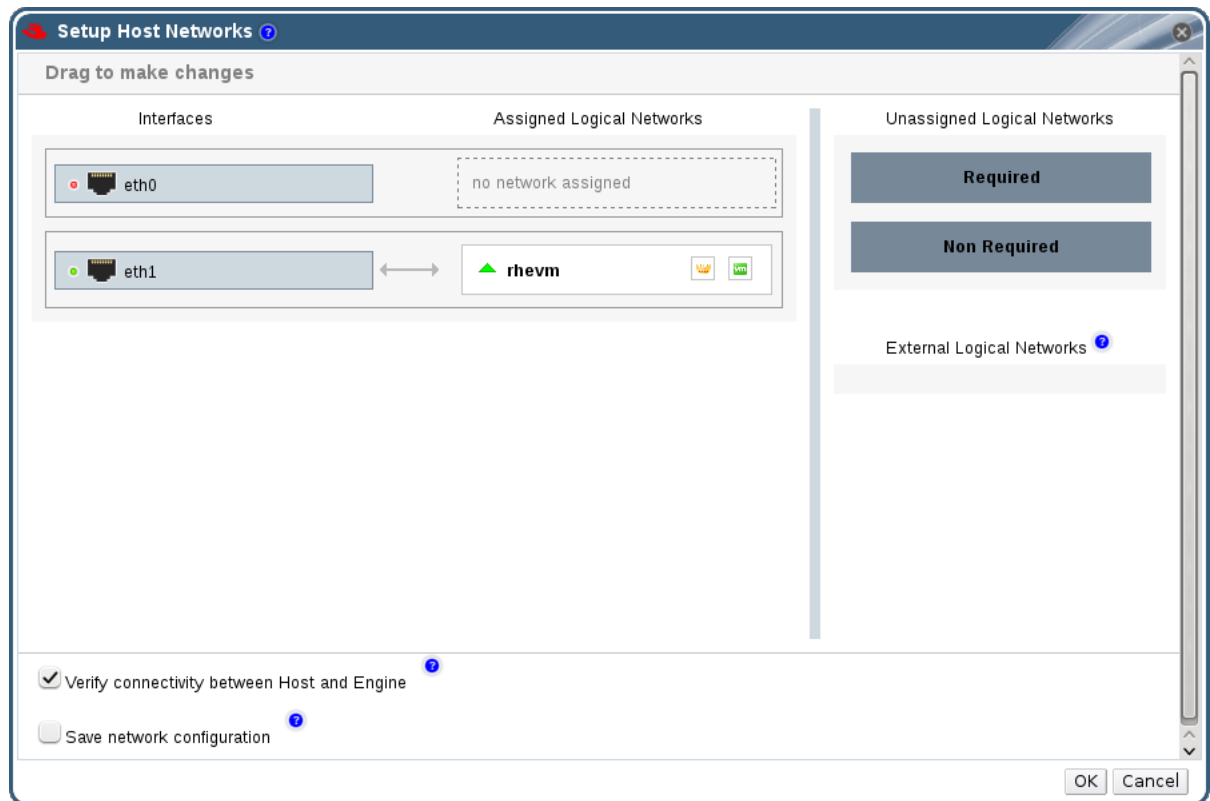


Figure 12.1. The **Setup Host Networks** window

4. Attach a logical network to a physical host network interface by selecting and dragging the logical network into the **Assigned Logical Networks** area next to the physical host network interface.

Alternatively, right-click the logical network and select a network interface from the drop-down menu.

5. Configure the logical network:
  - a. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.
  - b. Select a **Boot Protocol** from:
    - **None**,
    - **DHCP**, or
    - **Static**.

If you selected **Static**, enter the **IP**, **Subnet Mask**, and the **Gateway**.
  - c. Click **OK**.
  - d. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.
6. Select the **Verify connectivity between Host and Engine** check box to check network connectivity; this action will only work if the host is in maintenance mode.
7. Select the **Save network configuration** check box to make the changes persistent when the environment is rebooted.

8. Click **OK**.

## Result

You have assigned logical networks to and configured a physical host network interface.



## NOTE

If not all network interface cards for the host are displayed, click the **Refresh Capabilities** button to update the list of network interface cards available for that host.

[Report a bug](#)

## 12.4.2. Logical Network General Settings Explained

The table below describes the settings for the **General** tab of the **New Logical Network** and **Edit Logical Network** window.

Table 12.1. **New Logical Network** and **Edit Logical Network** Settings

Field Name	Description
<b>Name</b>	The name of the logical network. This text field has a 15-character limit and must be a unique name with any combination of uppercase and lowercase letters, numbers, hyphens, and underscores.
<b>Description</b>	The description of the logical network. This text field has a 40-character limit.
<b>Comment</b>	A field for adding plain text, human-readable comments regarding the logical network.
<b>Create on external provider</b>	Allows you to create the logical network to an OpenStack network service that has been added to the Manager as an external provider.  <b>External Provider</b> - Allows you to select the external provider on which the logical network will be created.
<b>Enable VLAN tagging</b>	VLAN tagging is a security feature that gives all network traffic carried on the logical network a special characteristic. VLAN-tagged traffic cannot be read by interfaces that do not also have that characteristic. Use of VLANs on logical networks also allows a single network interface to be associated with multiple, differently VLAN-tagged logical networks. Enter a numeric value in the text entry field if VLAN tagging is enabled.
<b>VM Network</b>	Select this option if only virtual machines use this network. If the network is used for traffic that does not involve virtual machines, such as storage communications, do not select this check box.

Field Name	Description
<b>Override MTU</b>	Set a custom maximum transmission unit for the logical network. You can use this to match the maximum transmission unit supported by your new logical network to the maximum transmission unit supported by the hardware it interfaces with. Enter a numeric value in the text entry field if <b>Override MTU</b> is selected.
<b>Network Label</b>	Allows you to specify a new label for the network or select from a existing labels already attached to host network interfaces. If you select an existing label, the logical network will be automatically assigned to all host network interfaces with that label.

[Report a bug](#)

### 12.4.3. Editing a Logical Network

#### Summary

Edit the settings of a logical network.

#### Procedure 12.3. Editing a Logical Network

1. Use the **Data Centers** resource tab, tree mode, or the search function to find and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.
3. Select a logical network and click **Edit** to open the **Edit Logical Network** window.
4. Edit the necessary settings.
5. Click **OK** to save the changes.

#### Result

You have updated the settings of your logical network.



#### NOTE

Multi-host network configuration is available on data centers with 3.1-or-higher compatibility, and automatically applies updated network settings to all of the hosts within the data center to which the network is assigned. Changes can only be applied when virtual machines using the network are down. You cannot rename a logical network that is already configured on a host. You cannot disable the **VM Network** option while virtual machines or templates using that network are running.

[Report a bug](#)



### 12.4.4. Explanation of Settings in the Manage Networks Window

The table below describes the settings for the **Manage Networks** window.

Table 12.2. Manage Networks Settings

Field	Description/Action
<b>Assign</b>	Assigns the logical network to all hosts in the cluster.
<b>Required</b>	A Network marked "required" must remain operational in order for the hosts associated with it to function properly. If a required network ceases to function, any hosts associated with it become non-operational.
<b>VM Network</b>	A logical network marked "VM Network" carries network traffic relevant to the virtual machine network.
<b>Display Network</b>	A logical network marked "Display Network" carries network traffic relevant to SPICE and to the virtual network controller.
<b>Migration Network</b>	A logical network marked "Migration Network" carries virtual machine traffic and storage migration traffic.

[Report a bug](#)

### 12.4.5. Adding Multiple VLANs to a Single Network Interface Using Logical Networks

#### Summary

Multiple VLANs can be added to a single network interface to separate traffic on the one host.



#### IMPORTANT

You must have created more than one logical network, all with the **Enable VLAN tagging** check box selected in the **New Logical Network** or **Edit Logical Network** windows.

#### Procedure 12.4. Adding Multiple VLANs to a Network Interface using Logical Networks

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select in the results list a host associated with the cluster to which your VLAN-tagged logical networks are assigned.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the data center.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Drag your VLAN-tagged logical networks into the **Assigned Logical Networks** area next to the physical network interface. The physical network interface can have multiple logical networks assigned due to the VLAN tagging.

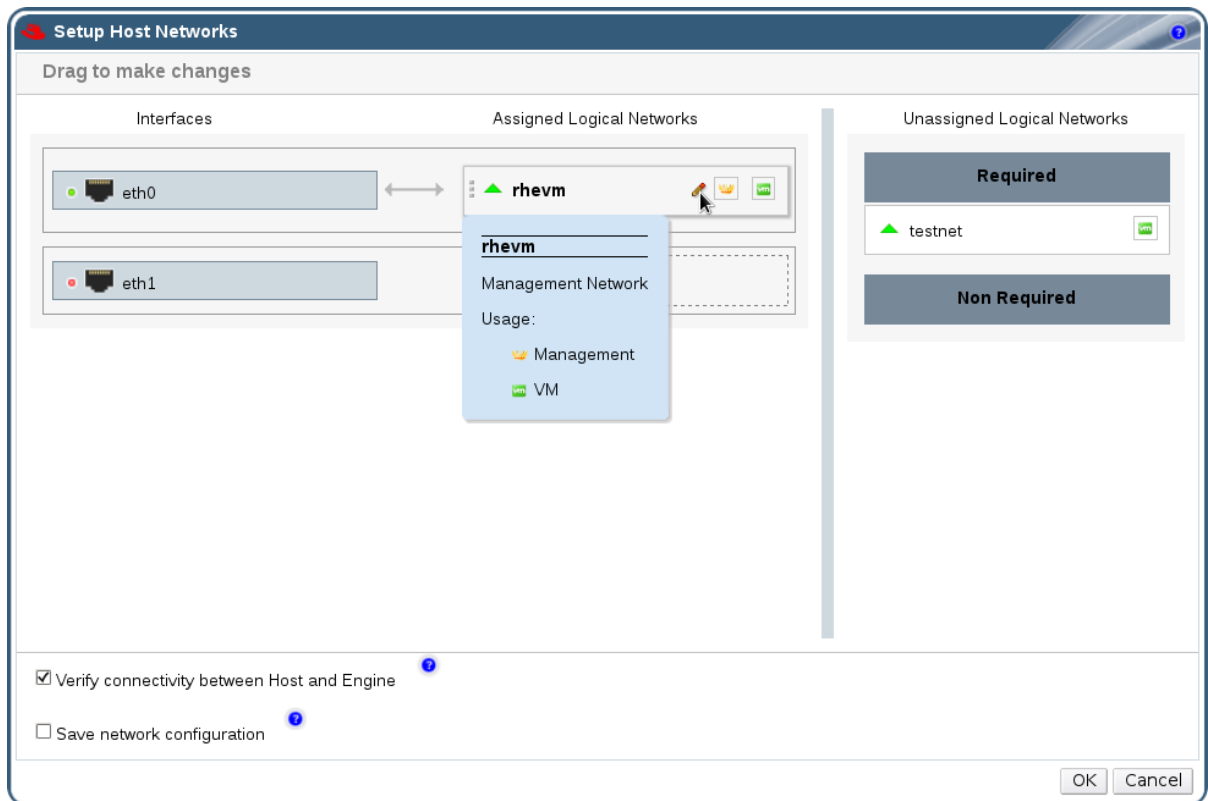


Figure 12.2. Setup Host Networks

5. Edit the logical networks by hovering your cursor over an assigned logical network and clicking the pencil icon to open the **Edit Network** window.

If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box.

Select a **Boot Protocol** from:

- **None**,
- **DHCP**, or
- **Static**,

Provide the **IP** and **Subnet Mask**.

Click **OK**.

6. Select the **Verify connectivity between Host and Engine** check box to run a network check; this will only work if the host is in maintenance mode.
7. Select the **Save network configuration** check box
8. Click **OK**.

Add the logical network to each host in the cluster by editing a NIC on each host in the cluster. After this is done, the network will become operational

## Result

You have added multiple VLAN-tagged logical networks to a single interface. This process can be repeated multiple times, selecting and editing the same network interface each time on each host to add logical networks with different VLAN tags to a single network interface.

[Report a bug](#)

## 12.4.6. Multiple Gateways

### Summary

Users can define the gateway, along with the IP address and subnet mask, for a logical network. This is necessary when multiple networks exist on a host and traffic should be routed through the specified network, rather than the default gateway.

If multiple networks exist on a host and the gateways are not defined, return traffic will be routed through the default gateway, which may not reach the intended destination. This would result in users being unable to ping the host.

Red Hat Enterprise Virtualization 3.4 handles multiple gateways automatically whenever an interface goes up or down.

### Procedure 12.5. Viewing or Editing the Gateway for a Logical Network

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the network interfaces attached to the host and their configurations.
3. Click the **Setup Host Networks** button to open the **Setup Host Networks** window.
4. Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.

### Result

The **Edit Management Network** window displays the network name, the boot protocol, and the IP, subnet mask, and gateway addresses. The address information can be manually edited by selecting a **Static** boot protocol.

[Report a bug](#)

## 12.4.7. Using the Networks Tab

The **Networks** resource tab provides a central location for users to perform network-related operations and search for networks based on each network's property or association with other resources.

All networks in the Red Hat Enterprise Virtualization environment display in the results list of the **Networks** tab. The **New**, **Edit** and **Remove** buttons allow you to create, change the properties of, and delete logical networks within data centers.

Click on each network name and use the **Clusters**, **Hosts**, **Virtual Machines**, **Templates**, and **Permissions** tabs in the details pane to perform functions including:

- Attaching or detaching the networks to clusters and hosts
- Removing network interfaces from virtual machines and templates

- Adding and removing permissions for users to access and manage networks

These functions are also accessible through each individual resource tab.

[Report a bug](#)

## 12.5. EXTERNAL PROVIDER NETWORKS

### 12.5.1. Importing Networks From External Providers

#### Summary

If an external provider offering networking services has been registered in the Manager, the networks provided by that provider can be imported into the Manager and used by virtual machines.

#### Procedure 12.6. Importing a Network From an External Provider

1. Click the **Networks** tab.
2. Click the **Import** button to open the **Import Networks** window.

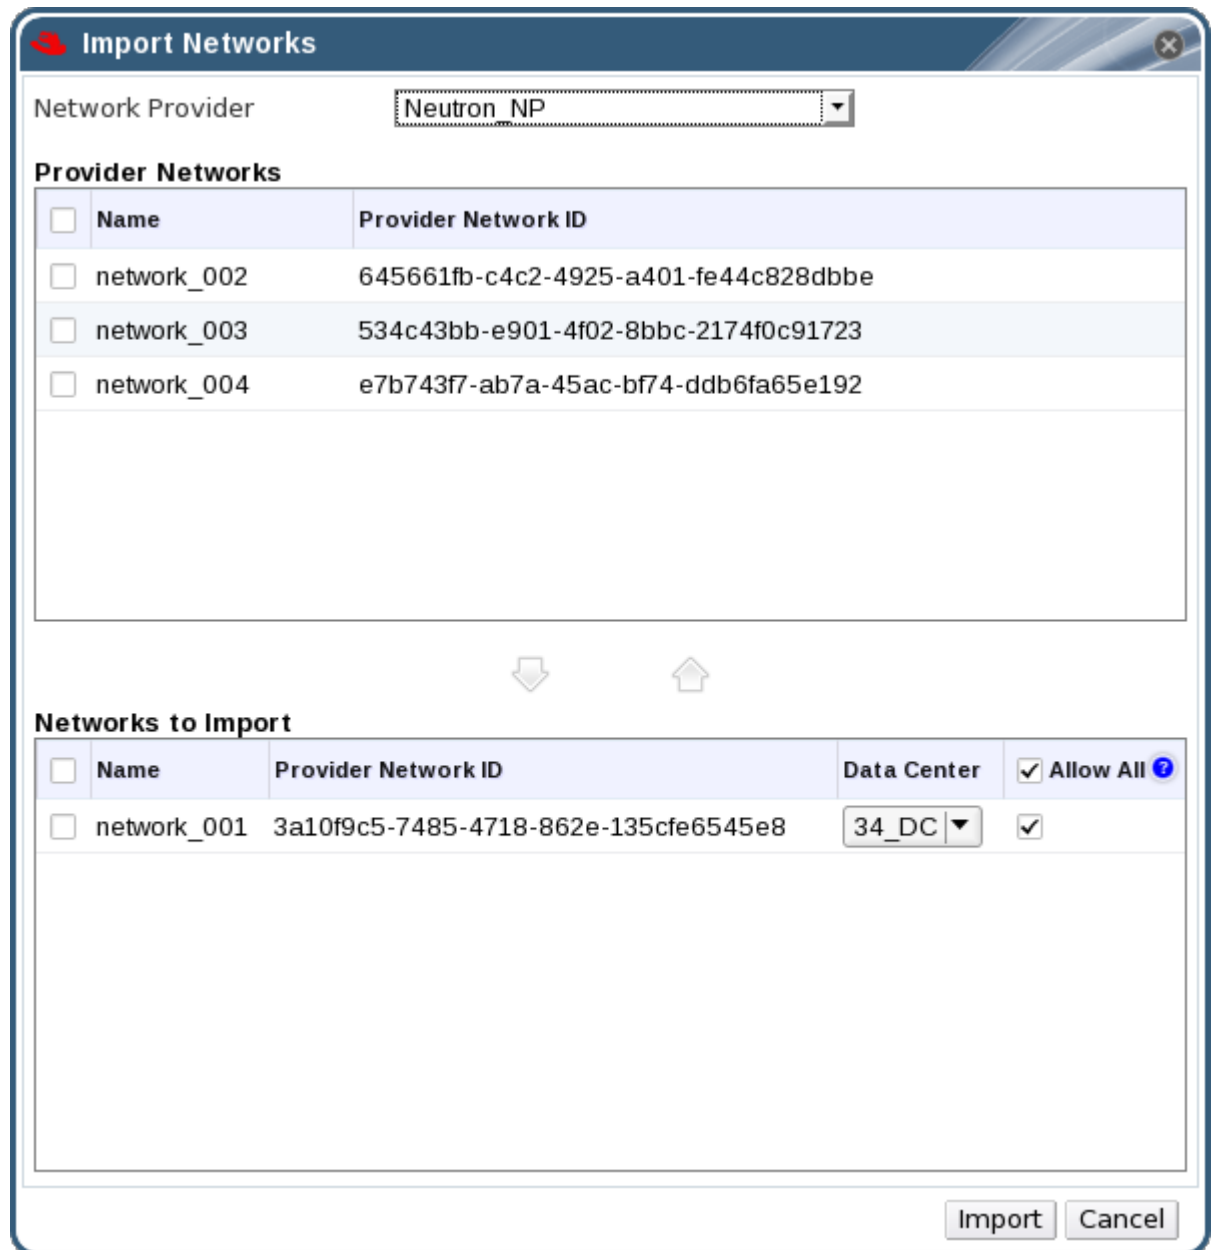


Figure 12.3. The Import Networks Window

- From the **Network Provider** drop-down list, select an external provider. The networks offered by that provider are automatically discovered and listed in the **Provider Networks** list.
- Using the check boxes, select the networks to import in the **Provider Networks** list and click the down arrow to move those networks into the **Networks to Import** list.
- From the **Data Center** drop-down list, select the data center into which the networks will be imported.
- Optionally, clear the **Allow All** check box for a network in the **Networks to Import** list to prevent that network from being available to all users.
- Click the **Import** button.

### Result

The selected networks are imported into the target data center and can now be used in the Manager.

**IMPORTANT**


External provider discovery and importing are Technology Preview features. 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.

[Report a bug](#)

### 12.5.2. Limitations to Using External Provider Networks

The following limitations apply to using logical networks imported from an external provider in a Red Hat Enterprise Virtualization environment.

- Logical networks offered by external providers must be used as virtual machine networks, and cannot be used as display networks.
- The same logical network can be imported more than once, but only to different data centers.
- You cannot edit logical networks offered by external providers in the Manager. To edit the details of a logical network offered by an external provider, you must edit the logical network directly from the OpenStack network service that provides that logical network.
- Port mirroring is not available for virtual network interface cards connected to logical networks offered by external providers.
- If a virtual machine uses a logical network offered by an external provider, that provider cannot be deleted from the Manager while the logical network is still in use by the virtual machine.
- Networks offered by external providers are non-required. As such, scheduling for clusters in which such logical networks have been imported will not take those logical networks into account during host selection. Moreover, it is the responsibility of the user to ensure the availability of the logical network on hosts in clusters in which such logical networks have been imported.

**IMPORTANT**

Logical networks imported from external providers are only compatible with Red Hat Linux hosts and cannot be assigned to virtual machines running on Red Hat Enterprise Virtualization Hypervisor hosts.

**IMPORTANT**

External provider discovery and importing are Technology Preview features. 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.

[Report a bug](#)

## 12.5.3. Configuring Subnets on External Provider Logical Networks

### 12.5.3.1. Configuring Subnets on External Provider Logical Networks

A logical network provided by an external provider can only assign IP addresses to virtual machines if one or more subnets have been defined on that logical network. If no subnets are defined, virtual machines will not be assigned IP addresses. If there is one subnet, virtual machines will be assigned an IP address from that subnet, and if there are multiple subnets, virtual machines will be assigned an IP address from any of the available subnets. The DHCP service provided by the Neutron instance on which the logical network is hosted is responsible for assigning these IP addresses.

While the Red Hat Enterprise Virtualization Manager automatically discovers predefined subnets on imported logical networks, you can also add or remove subnets to or from logical networks from within the Manager.

[Report a bug](#)

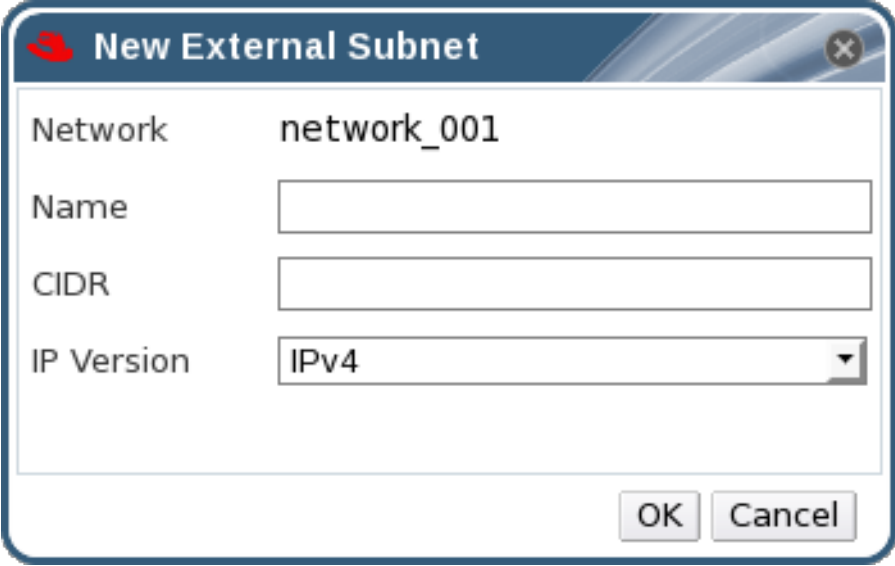
### 12.5.3.2. Adding Subnets to External Provider Logical Networks

#### Summary

Create a subnet on a logical network provided by an external provider

#### Procedure 12.7. Adding Subnets to External Provider Logical Networks

1. Click the **Networks** tab.
2. Click the logical network provided by an external provider to which the subnet will be added.
3. Click the **Subnets** tab in the details pane.
4. Click the **New** button to open the **New External Subnet** window.



Network	network_001
Name	<input type="text"/>
CIDR	<input type="text"/>
IP Version	IPv4

OK Cancel

Figure 12.4. The New External Subnet Window

5. Enter a **Name** and **CIDR** for the new subnet.
6. From the **IP Version** drop-down menu, select either **IPv4** or **IPv6**.

7. Click **OK**.

### Result

A new subnet is created on the logical network.

[Report a bug](#)

### 12.5.3.3. Removing Subnets from External Provider Logical Networks

#### Summary

Remove a subnet from a logical network provided by an external provider

#### Procedure 12.8. Removing Subnets from External Provider Logical Networks

1. Click the **Networks** tab.
2. Click the logical network provided by an external provider from which the subnet will be removed.
3. Click the **Subnets** tab in the details pane.
4. Click the subnet to remove.
5. Click the **Remove** button and click **OK** when prompted.

### Result

The subnet is removed from the logical network.

[Report a bug](#)

## 12.6. BONDING

### 12.6.1. Bonding Logic in Red Hat Enterprise Virtualization

The Red Hat Enterprise Virtualization Manager Administration Portal allows you to create bond devices using a graphical interface. There are several distinct bond creation scenarios, each with its own logic.

Two factors that affect bonding logic are:

- Are either of the devices already carrying logical networks?
- Are the devices carrying compatible logical networks? A single device cannot carry both VLAN tagged and non-VLAN tagged logical networks.

**Table 12.3. Bonding Scenarios and Their Results**

Bonding Scenario	Result
------------------	--------



Bonding Scenario	Result
NIC + NIC	<p>The <b>Create New Bond</b> window is displayed, and you can configure a new bond device.</p> <p>If the network interfaces carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>
NIC + Bond	<p>The NIC is added to the bond device. Logical networks carried by the NIC and the bond are all added to the resultant bond device if they are compatible.</p> <p>If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>
Bond + Bond	<p>If the bond devices are not attached to logical networks, or are attached to compatible logical networks, a new bond device is created. It contains all of the network interfaces, and carries all logical networks, of the component bond devices. The <b>Create New Bond</b> window is displayed, allowing you to configure your new bond.</p> <p>If the bond devices carry incompatible logical networks, the bonding operation fails until you detach incompatible logical networks from the devices forming your new bond.</p>

[Report a bug](#)

## 12.6.2. Bonding Modes

Red Hat Enterprise Virtualization supports the following common bonding modes:

- Mode 1 (active-backup policy) sets all interfaces to the backup state while one remains active. Upon failure on the active interface, a backup interface replaces it as the only active interface in the bond. The MAC address of the bond in mode 1 is visible on only one port (the network adapter), to prevent confusion for the switch. Mode 1 provides fault tolerance and is supported in Red Hat Enterprise Virtualization.
- Mode 2 (XOR policy) selects an interface to transmit packages to based on the result of an XOR operation on the source and destination MAC addresses modulo NIC slave count. This calculation ensures that the same interface is selected for each destination MAC address used. Mode 2 provides fault tolerance and load balancing and is supported in Red Hat Enterprise Virtualization.
- Mode 4 (IEEE 802.3ad policy) creates aggregation groups for which included interfaces share the speed and duplex settings. Mode 4 uses all interfaces in the active aggregation group in accordance with the IEEE 802.3ad specification and is supported in Red Hat Enterprise Virtualization.
- Mode 5 (adaptive transmit load balancing policy) ensures the outgoing traffic distribution is

according to the load on each interface and that the current interface receives all incoming traffic. If the interface assigned to receive traffic fails, another interface is assigned the receiving role instead. Mode 5 is supported in Red Hat Enterprise Virtualization.

[Report a bug](#)

### 12.6.3. Creating a Bond Device Using the Administration Portal

#### Summary

You can bond compatible network devices together. This type of configuration can increase available bandwidth and reliability. You can bond multiple network interfaces, pre-existing bond devices, and combinations of the two.

A bond cannot carry both vlan tagged and non-vlan traffic.

#### Procedure 12.9. Creating a Bond Device using the Administration Portal

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click the **Network Interfaces** tab in the details pane to list the physical network interfaces attached to the host.
3. Click **Setup Host Networks** to open the **Setup Host Networks** window.
4. Select and drag one of the devices over the top of another device and drop it to open the **Create New Bond** window. Alternatively, right-click the device and select another device from the drop-down menu.

If the devices are incompatible, for example one is vlan tagged and the other is not, the bond operation fails with a suggestion on how to correct the compatibility issue.

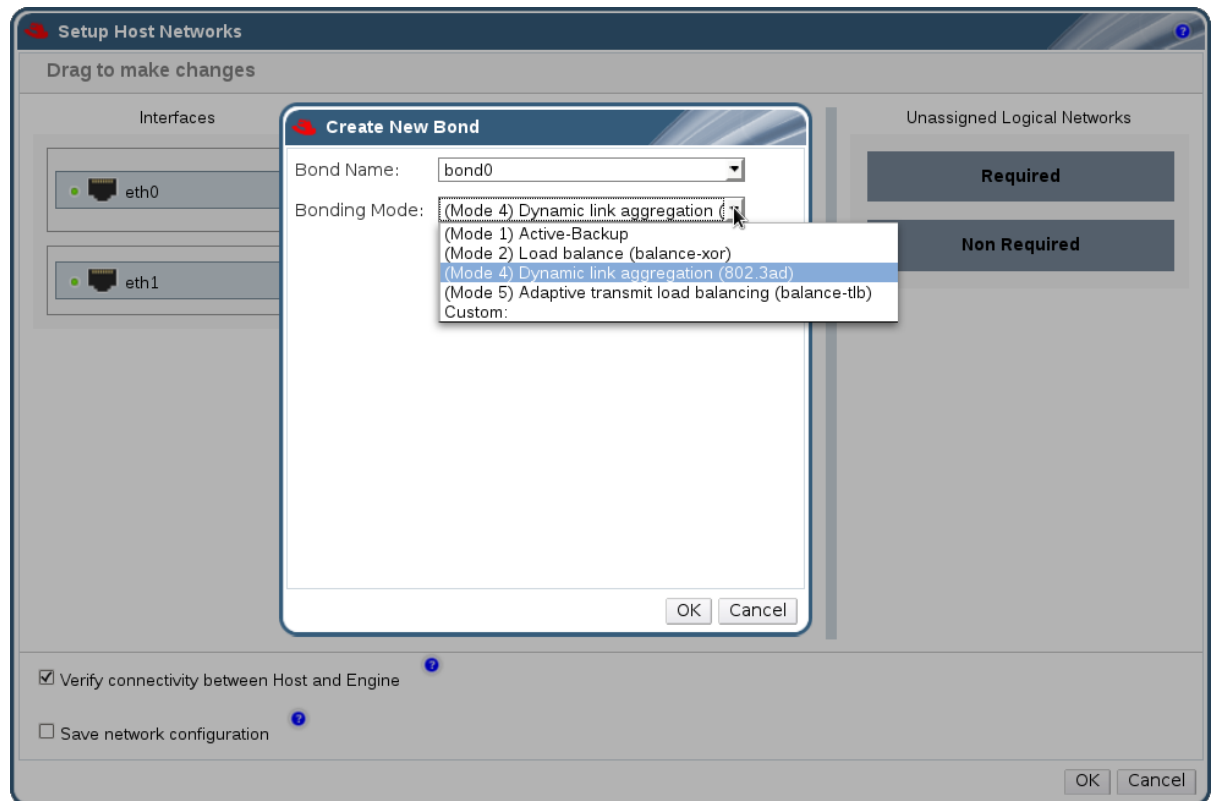


Figure 12.5. Bond Devices Window

5. Select the **Bond Name** and **Bonding Mode** from the drop-down menus.

Bonding modes 1, 2, 4, and 5 can be selected. Any other mode can be configured using the **Custom** option.

6. Click **OK** to create the bond and close the **Create New Bond** window.
7. Assign a logical network to the newly created bond device.
8. Optionally choose to **Verify connectivity between Host and Engine** and **Save network configuration**.
9. Click **OK** accept the changes and close the **Setup Host Networks** window.

#### Result:

Your network devices are linked into a bond device and can be edited as a single interface. The bond device is listed in the **Network Interfaces** tab of the details pane for the selected host.

Bonding must be enabled for the ports of the switch used by the host. The process by which bonding is enabled is slightly different for each switch; consult the manual provided by your switch vendor for detailed information on how to enable bonding.

[Report a bug](#)

#### 12.6.4. Example Uses of Custom Bonding Options with Host Interfaces

You can create customized bond devices by selecting **Custom** from the **Bonding Mode** of the **Create New Bond** window. The following examples should be adapted for your needs. For a comprehensive list of bonding options and their descriptions, see the [Linux Ethernet Bonding Driver HOWTO](#) on Kernel.org.

### Example 12.1. xmit\_hash\_policy

This option defines the transmit load balancing policy for bonding modes 2 and 4. For example, if the majority of your traffic is between many different IP addresses, you may want to set a policy to balance by IP address. You can set this load-balancing policy by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=4 xmit_hash_policy=layer2+3
```

### Example 12.2. ARP Monitoring

ARP monitor is useful for systems which can't or don't report link-state properly via ethtool. Set an **arp\_interval** on the bond device of the host by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1 arp_interval=1 arp_ip_target=192.168.0.2
```

### Example 12.3. Primary

You may want to designate a NIC with higher throughput as the primary interface in a bond device. Designate which NIC is primary by selecting a **Custom** bonding mode, and entering the following into the text field:

```
mode=1 primary=eth0
```

[Report a bug](#)

## 12.7. REMOVING LOGICAL NETWORKS

### 12.7.1. Removing a Logical Network

#### Summary

Remove a logical network from the Manager.

#### Procedure 12.10. Removing Logical Networks

1. Use the **Data Centers** resource tab, tree mode, or the search function to find and select the data center of the logical network in the results list.
2. Click the **Logical Networks** tab in the details pane to list the logical networks in the data center.
3. Select a logical network and click **Remove** to open the **Remove Logical Network(s)** window.
4. Optionally, select the **Remove external network(s) from the provider(s) as well** check box to remove the logical network both from the Manager and from the external provider if the network is provided by an external provider.

5. Click **OK**.

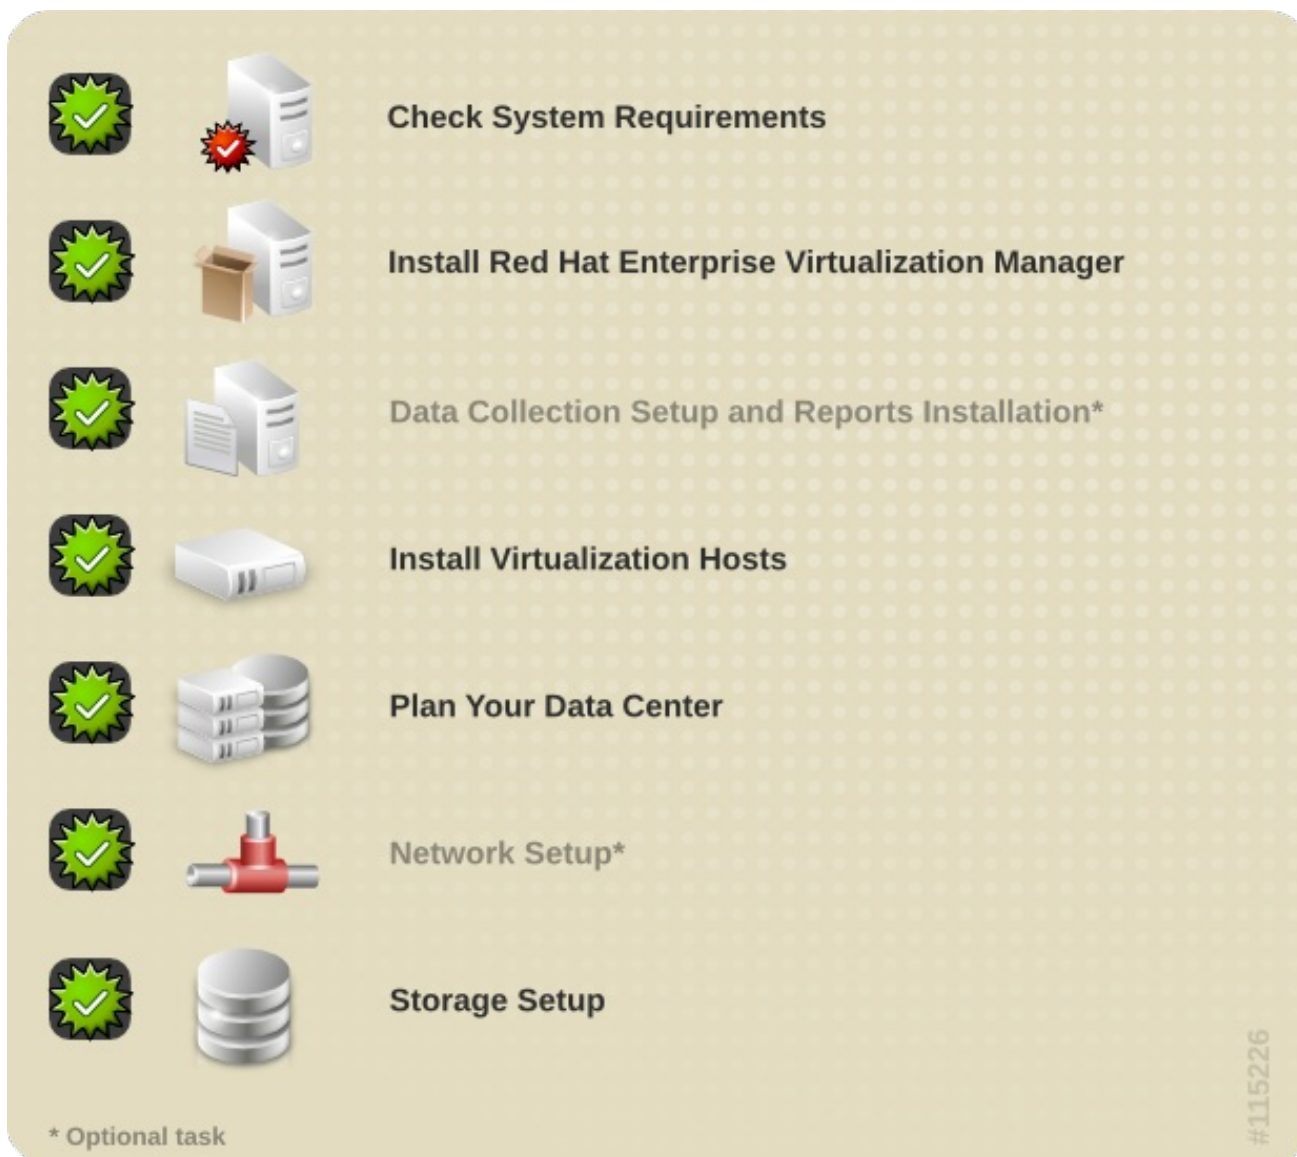
**Result**

The logical network is removed from the Manager and is no longer available. If the logical network was provided by an external provider and you elected to remove the logical network from that external provider, it is removed from the external provider and is no longer available on that external provider as well.

[Report a bug](#)

## CHAPTER 13. CONFIGURING STORAGE

### 13.1. WORKFLOW PROGRESS - STORAGE SETUP



[Report a bug](#)

### 13.2. INTRODUCTION TO STORAGE IN RED HAT ENTERPRISE VIRTUALIZATION

Red Hat Enterprise Virtualization uses a centralized storage system for virtual machine disk images, ISO files and snapshots. Storage networking can be implemented using:

- Network File System (NFS)
- GlusterFS exports
- Other POSIX compliant file systems
- Internet Small Computer System Interface (iSCSI)
- Local storage attached directly to the virtualization hosts

- Fibre Channel Protocol (FCP)
- Parallel NFS (pNFS)

Setting up storage is a prerequisite for a new data center because a data center cannot be initialized unless storage domains are attached and activated.

As a Red Hat Enterprise Virtualization system administrator, you need to create, configure, attach and maintain storage for the virtualized enterprise. You should be familiar with the storage types and their use. Read your storage array vendor's guides, and refer to the *Red Hat Enterprise Linux Storage Administration Guide* for more information on the concepts, protocols, requirements or general usage of storage.

The Red Hat Enterprise Virtualization platform enables you to assign and manage storage using the Administration Portal's **Storage** tab. The **Storage** results list displays all the storage domains, and the details pane shows general information about the domain.

Red Hat Enterprise Virtualization platform has three types of storage domains:

- **Data Domain:** A data domain holds the virtual hard disks and OVF files of all the virtual machines and templates in a data center. In addition, snapshots of the virtual machines are also stored in the data domain.

The data domain cannot be shared across data centers. Storage 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.

You must attach a data domain to a data center before you can attach domains of other types to it.


- **ISO Domain:** ISO domains store ISO files (or logical CDs) used to install and boot operating systems and applications for the virtual machines. An ISO domain removes the data center's need for physical media. An ISO domain can be shared across different data centers.
- **Export Domain:** Export domains are temporary storage repositories that are used to copy and move images between data centers and Red Hat Enterprise 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.



### IMPORTANT

Support for export storage domains backed by storage on anything other than NFS is being deprecated. While existing export storage domains imported from Red Hat Enterprise Virtualization 2.2 environments remain supported new export storage domains must be created on NFS storage.

Only commence configuring and attaching storage for your Red Hat Enterprise Virtualization environment once you have determined the storage needs of your data center(s).



### IMPORTANT

To add storage domains you must be able to successfully access the Administration Portal, and there must be at least one host connected with a status of **Up**.

[Report a bug](#)

## 13.3. PREPARING NFS STORAGE

### Summary

These steps must be taken to prepare an NFS file share on a server running Red Hat Enterprise Linux 6 for use with Red Hat Enterprise Virtualization.

### Procedure 13.1. Preparing NFS Storage

#### 1. Install `nfs-utils`

NFS functionality is provided by the `nfs-utils` package. Before file shares can be created, check that the package is installed by querying the RPM database for the system:

```
$ rpm -qi nfs-utils
```

If the `nfs-utils` package is installed then the package information will be displayed. If no output is displayed then the package is not currently installed. Install it using `yum` while logged in as the **root** user:

```
# yum install nfs-utils
```

#### 2. Configure Boot Scripts

To ensure that NFS shares are always available when the system is operational both the **nfs** and **rpcbind** services must start at boot time. Use the **chkconfig** command while logged in as **root** to modify the boot scripts.

```
# chkconfig --add rpcbind
# chkconfig --add nfs
# chkconfig rpcbind on
# chkconfig nfs on
```

Once the boot script configuration has been done, start the services for the first time.

```
# service rpcbind start
# service nfs start
```

#### 3. Create Directory

Create the directory you wish to share using NFS.

```
# mkdir /exports/iso
```

Replace `/exports/iso` with the name, and path of the directory you wish to use.

#### 4. Export Directory

To be accessible over the network using NFS the directory must be exported. NFS exports are controlled using the `/etc/exports` configuration file. Each export path appears on a separate line followed by a tab character and any additional NFS options. Exports to be attached to the Red Hat Enterprise Virtualization Manager must have the `read`, and `write`, options set.

To grant read, and write access to `/exports/iso` using NFS for example you add the following line to the `/etc/exports` file.

```
/exports/iso *(rw)
```



Again, replace `/exports/iso` with the name, and path of the directory you wish to use.

### 5. Reload NFS Configuration

For the changes to the `/etc/exports` file to take effect the service must be told to reload the configuration. To force the service to reload the configuration run the following command as **root**:

```
# service nfs reload
```

### 6. Set Permissions

The NFS export directory must be configured for read write access and must be owned by `vdsm:kvm`. If these users do not exist on your external NFS server use the following command, assuming that `/exports/iso` is the directory to be used as an NFS share.

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

The permissions on the directory must be set to allow read and write access to both the owner and the group. The owner should also have execute access to the directory. The permissions are set using the **chmod** command. The following command arguments set the required permissions on the `/exports/iso` directory.

```
# chmod 0755 /exports/iso
```

### Result

The NFS file share has been created, and is ready to be attached by the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

## 13.4. ATTACHING NFS STORAGE

### Summary

An NFS type **Storage Domain** is a mounted NFS share that is attached to a data center. It is used to provide storage for virtualized guest images and ISO boot media. Once NFS storage has been exported it must be attached to the Red Hat Enterprise Virtualization Manager using the Administration Portal.

NFS data domains can be added to NFS data centers. You can add NFS, ISO, and export storage domains to data centers of any type.

### Procedure 13.2. Attaching NFS Storage

1. Click the **Storage** resource tab to list the existing storage domains.
2. Click **New Domain** to open the **New Domain** window.

The screenshot shows the 'New Domain' dialog box with the following fields and options:

- Name:** Text input field.
- Description:** Text input field.
- Data Center:** Drop-down menu with 'Default (NFS)' selected.
- Comment:** Text input field.
- Domain Function / Storage Type:** Drop-down menu with 'Data / NFS' selected.
- Format:** Drop-down menu with 'V3' selected.
- Use Host:** Drop-down menu.
- Export Path:** Text input field with a tooltip: *Remote path to NFS export, takes either the form: FQDN:/path or IP:/path e.g. server.example.com:/export/VMs*
- Advanced Parameters:**
  - Override Default Options
  - NFS Version:** Drop-down menu with 'V3 (default)' selected.
  - Retransmissions (#):** Text input field.
  - Timeout (deciseconds):** Text input field.

At the bottom right, there are 'OK' and 'Cancel' buttons.

Figure 13.1. NFS Storage


3. Enter the **Name** of the storage domain.
4. Select the **Data Center**, **Domain Function / Storage Type**, and **Use Host** from the drop-down menus.

If applicable, select the **Format** from the drop-down menu.

5. Enter the **Export Path** to be used for the storage domain.

The export path should be in the format of **192.168.0.10:/data** or **domain.example.com:/data**

6. Click **Advanced Parameters** to enable further configurable settings. It is recommended that the values of these parameters not be modified.



### IMPORTANT

All communication to the storage domain is from the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must be attached to the chosen Data Center before the storage is configured.

7. Click **OK** to create the storage domain and close the window.

### Result

The new NFS data domain is displayed on the **Storage** tab with a status of **Locked** while the disk prepares. It is automatically attached to the data center upon completion.

[Report a bug](#)

## 13.5. PREPARING PNFS STORAGE

Support for Parallel NFS (pNFS) as part of the NFS v4.1 standard is available as of Red Hat Enterprise Linux 6.4. The pNFS architecture improves the scalability of NFS, with possible improvements to performance. That is, when a server implements pNFS as well, a client is able to access data through multiple servers concurrently. The pNFS protocol supports three storage protocols or layouts: files, objects, and blocks. Red Hat Enterprise Linux 6.4 supports only the "files" layout type.

To enable support for pNFS functionality, use one of the following mount options on mounts from a pNFS-enabled server:

```
-o minorversion=1
```

or

```
-o v4.1
```

Set the permissions of the pNFS path so that Red Hat Enterprise Virtualization can access them:

```
# chown 36:36 [path to pNFS resource]
```

After the server is pNFS-enabled, the `nfs_layout_nfsv41_files` kernel is automatically loaded on the first mount. Verify that the module was loaded:

```
$ lsmod | grep nfs_layout_nfsv41_files
```

Another way to verify a successful NFSv4.1 mount is with the `mount` command. The mount entry in the output should contain `minorversion=1`.

[Report a bug](#)

## 13.6. ATTACHING PNFS STORAGE

### Summary

A pNFS type **Storage Domain** is a mounted pNFS share attached to a data center. It provides storage for virtualized guest images and ISO boot media. After you have exported pNFS storage, it must be attached to the Red Hat Enterprise Virtualization Manager using the Administration Portal.

### Procedure 13.3. Attaching pNFS Storage

1. Click the **Storage** resource tab to list the existing storage domains.
2. Click **New Domain** to open the **New Domain** window.

**New Domain**

Name:

Data Center:

Domain Function / Storage Type:  Format:

Use Host:

Path:   
Path to device to mount / remote export

VFS Type:

Mount Options:

OK Cancel

**Figure 13.2. NFS Storage**


3. Enter the **Name** of the storage domain.
4. Select the **Data Center**, **Domain Function / Storage Type**, and **Use Host** from the drop-down menus.

If applicable, select the **Format** from the drop-down menu.

5. Enter the **Export Path** to be used for the storage domain.

The export path should be in the format of **192.168.0.10:/data** or **domain.example.com:/data**

6. In the **VFS Type** field, enter **nfs4**.
7. In the **Mount Options** field, enter **minorversion=1**.



### IMPORTANT

All communication to the storage domain comes from the selected host and not from the Red Hat Enterprise Virtualization Manager. At least one active host must be attached to the chosen Data Center before the storage is configured.

8. Click **OK** to create the storage domain and close the window.

### Result

The new pNFS data domain is displayed on the **Storage** tab with a status of **Locked** while the disk prepares. It is automatically attached to the data center upon completion.

[Report a bug](#)

## 13.7. ADDING ISCSI STORAGE

### Summary

Red Hat Enterprise Virtualization platform supports iSCSI 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.

For information regarding the setup and configuration of iSCSI on Red Hat Enterprise Linux, see the *Red Hat Enterprise Linux Storage Administration Guide* .



### NOTE

You can only add an iSCSI storage domain to a data center that is set up for iSCSI storage type.

### Procedure 13.4. Adding iSCSI Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click the **New Domain** button to open the **New Domain** window.
3. Enter the **Name** of the new storage domain.

**New Domain**

Name: ISCSI-data  
Data Center: (none)

Domain Function / Storage Type: Data / iSCSI  
Format: V3

Use Host: rhev4

Discover Targets

Address:   
Port: 3260  
Discover

User Authentication:  
CHAP username:  CHAP password:

Login All

Target Name	Address	Port


OK Cancel

Figure 13.3. New iSCSI Domain

4. Use the **Data Center** drop-down menu to select an iSCSI data center.

If you do not yet have an appropriate iSCSI data center, select **(none)**.

5. Use the drop-down menus to select the **Domain Function / Storage Type** and the **Format**. 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 via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is configured.

7. The Red Hat Enterprise Virtualization Manager is able to map either iSCSI targets to LUNs, or LUNs to iSCSI targets. The **New Domain** window automatically displays known targets with

unused LUNs when iSCSI is selected as the storage type. If the target that you are adding storage from is not listed then you can use target discovery to find it, otherwise proceed to the next step.

### iSCSI Target Discovery

1. Click **Discover Targets** to enable target discovery options. When targets have been discovered and logged in to, the **New Domain** window automatically displays targets with LUNs unused by the environment.



#### NOTE

LUNs used externally to the environment are also displayed.

You can use the **Discover Targets** options to add LUNs on many targets, or multiple paths to the same LUNs.

2. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.
3. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.
4. If the Challenge Handshake Authentication Protocol (CHAP) is being used to secure the storage, select the **User Authentication** check box. Enter the **CHAP user name** and **CHAP password**.
5. Click the **Discover** button.
6. Select the target to use from the discovery results and click the **Login** button.

Alternatively, click the **Login All** to log in to all of the discovered targets.

8. Click the **+** button next to the desired target. This will expand the entry and display all unused LUNs attached to the target.
9. Select the check box for each LUN that you are using to create the storage domain.
10. Click **OK** to create the storage domain and close the window.

#### Result

The new iSCSI storage domain displays on the storage tab. This can take up to 5 minutes.

[Report a bug](#)

## 13.8. ADDING FCP STORAGE

### Summary

Red Hat Enterprise 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 Enterprise 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, please refer to the *Storage Administration Guide* and *DM Multipath Guide*.



## NOTE

You can only add an FCP storage domain to a data center that is set up for FCP storage type.

### Procedure 13.5. Adding FCP Storage

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

**New Domain**

Name:

Data Center:

Domain Function / Storage Type:  Format:

Use Host:

<input type="checkbox"/>	LUN ID	Dev. Size	#path	In Use	Vendor ID	Product ID	Serial


OK Cancel

Figure 13.4. 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 / Storage Type** and the **Format**. 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 via the selected host and not directly from the Red Hat Enterprise Virtualization Manager. At least one active host must exist in the system, and be attached to the chosen data center, before the storage is 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. Click **OK** to create the storage domain and close the window.

### Result

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.

[Report a bug](#)

## 13.9. PREPARING LOCAL STORAGE

### Summary

A local storage domain can be set up on a host. When you set up host to use local storage, the host automatically gets added to a new data center and cluster that no other hosts can be added to. Multiple host clusters require that all hosts have access to all storage domains, which is not possible with local storage. Virtual machines created in a single host cluster cannot be migrated, fenced or scheduled.



### IMPORTANT

On Red Hat Enterprise Virtualization Hypervisors the only path permitted for use as local storage is **/data/images**. This directory already exists with the correct permissions on Hypervisor installations. The steps in this procedure are only required when preparing local storage on Red Hat Enterprise Linux virtualization hosts.

### Procedure 13.6. Preparing Local Storage

1. On the virtualization host, create the directory to be used for the local storage.

```
# mkdir -p /data/images
```

2. Ensure that the directory has permissions allowing read/write access to the **vds** user (UID 36) and **kvm** group (GID 36).

```
# chown 36:36 /data /data/images
```

```
# chmod 0755 /data /data/images
```

### Result

Your local storage is ready to be added to the Red Hat Enterprise Virtualization environment.

[Report a bug](#)

## 13.10. ADDING LOCAL STORAGE

### Summary

Storage local to your host has been prepared. Now use the Manager to add it to the host.

Adding local storage to a host in this manner causes the host to be put in a new data center and cluster. The local storage configuration window combines the creation of a data center, a cluster, and storage into a single process.

### Procedure 13.7. Adding Local Storage

1. Use the **Hosts** resource tab, tree mode, or the search function to find and select the host in the results list.
2. Click **Maintenance** to open the **Maintenance Host(s)** confirmation window.
3. Click **OK** to initiate maintenance mode.
4. Click **Configure Local Storage** to open the **Configure Local Storage** window.

**Configure Local Storage**

**General**

Memory Optimization

**Data Center** Name  [Edit](#)

**Cluster** Name  [Edit](#)

**Storage** Name  [Edit](#)

Set the path to your local storage:

OK Cancel

**Figure 13.5. Configure Local Storage Window**

5. Click the **Edit** buttons next to the **Data Center**, **Cluster**, and **Storage** fields to configure and name the local storage domain.
6. Set the path to your local storage in the text entry field.

7. If applicable, select the **Memory Optimization** tab to configure the memory optimization policy for the new local storage cluster.
8. Click **OK** to save the settings and close the window.

### Result


Your host comes online in a data center of its own.

[Report a bug](#)

## 13.11. POSIX COMPLIANT FILE SYSTEM STORAGE IN RED HAT ENTERPRISE VIRTUALIZATION

Red Hat Enterprise Virtualization 3.1 and higher supports the use of POSIX (native) file systems for storage. POSIX file system support allows you to mount file systems using the same mount options that you would normally use when mounting them manually from the command line. This functionality is intended to allow access to storage not exposed using NFS, iSCSI, or FCP.

Any POSIX compliant filesystem used as a storage domain in Red Hat Enterprise Virtualization **MUST** support sparse files and direct I/O. The Common Internet File System (CIFS), for example, does not support direct I/O, making it incompatible with Red Hat Enterprise Virtualization.



### IMPORTANT

Do *not* mount NFS storage by creating a POSIX compliant file system Storage Domain. Always create an NFS Storage Domain instead.

[Report a bug](#)

## 13.12. ATTACHING POSIX COMPLIANT FILE SYSTEM STORAGE

### Summary

You want to use a POSIX compliant file system that is not exposed using NFS, iSCSI, or FCP as a storage domain.

### Procedure 13.8. Attaching POSIX Compliant File System Storage

1. Click the **Storage** resource tab to list the existing storage domains in the results list.
2. Click **New Domain** to open the **New Domain** window.

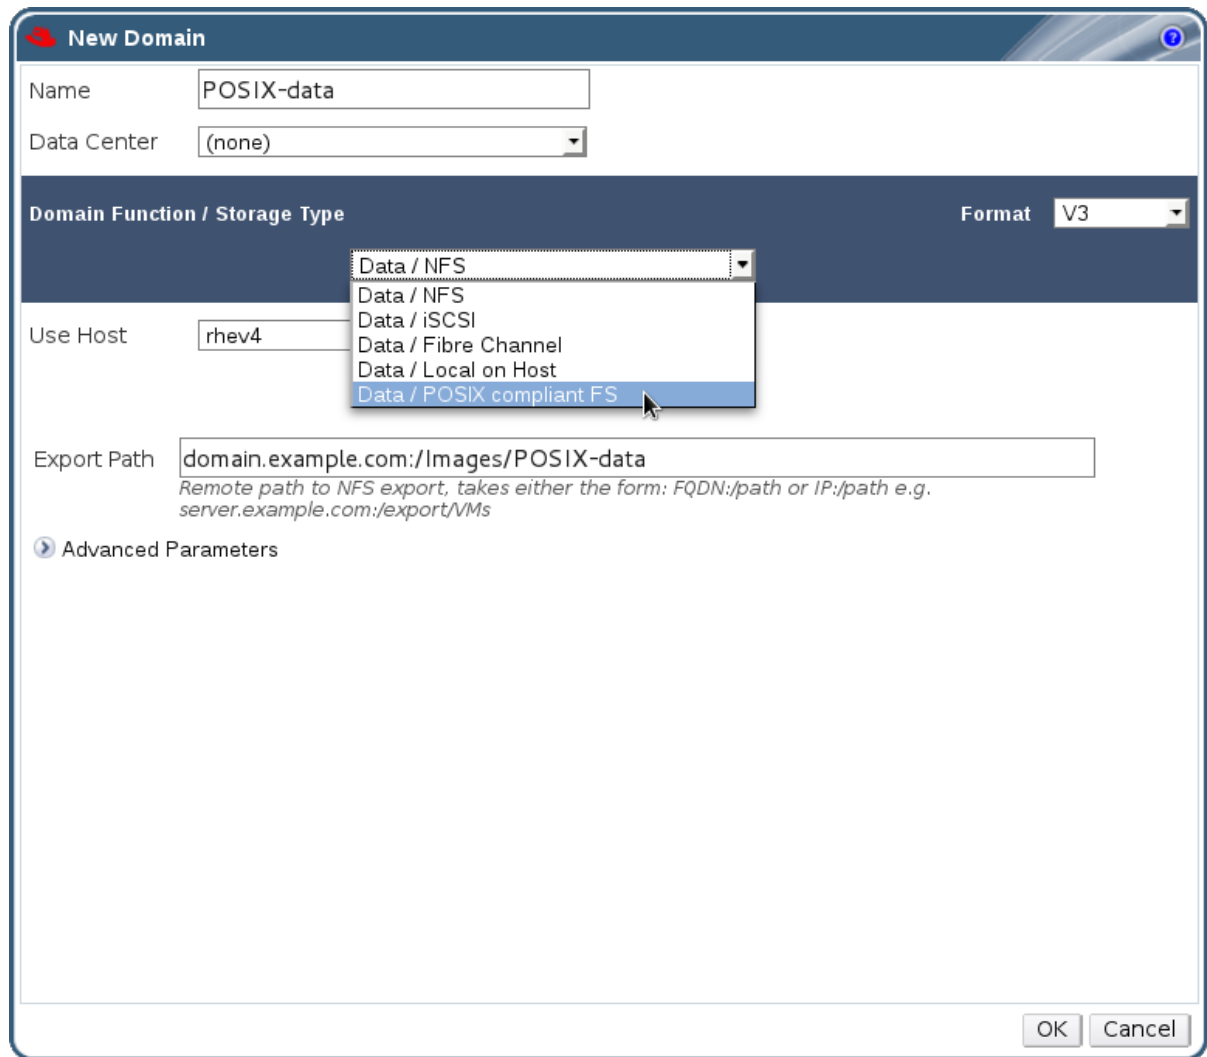


Figure 13.6. POSIX Storage

3. Enter the **Name** for the storage domain.
4. Select the **Data Center** to be associated with the storage domain. The Data Center selected must be of type **POSIX (POSIX compliant FS)**. Alternatively, select **(none)**.
5. Select **Data / POSIX compliant FS** from the **Domain Function / Storage Type** drop-down menu.

If applicable, select the **Format** from the drop-down menu.

6. Select a host from the **Use Host** drop-down menu. Only hosts within the selected data center will be listed. The host that you select will be used to connect the storage domain.
7. Enter the **Path** to the POSIX file system, as you would normally provide it to the **mount** command.
8. Enter the **VFS Type**, as you would normally provide it to the **mount** command using the **-t** argument. See **man mount** for a list of valid VFS types.
9. Enter additional **Mount Options**, as you would normally provide them to the **mount** command using the **-o** argument. The mount options should be provided in a comma-separated list. See **man mount** for a list of valid mount options.
10. Click **OK** to attach the new Storage Domain and close the window.

## Result

You have used a supported mechanism to attach an unsupported file system as a storage domain.

[Report a bug](#)

## 13.13. ENABLING GLUSTER PROCESSES ON RED HAT STORAGE NODES

### Summary

This procedure explains how to allow Gluster processes on Red Hat 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.

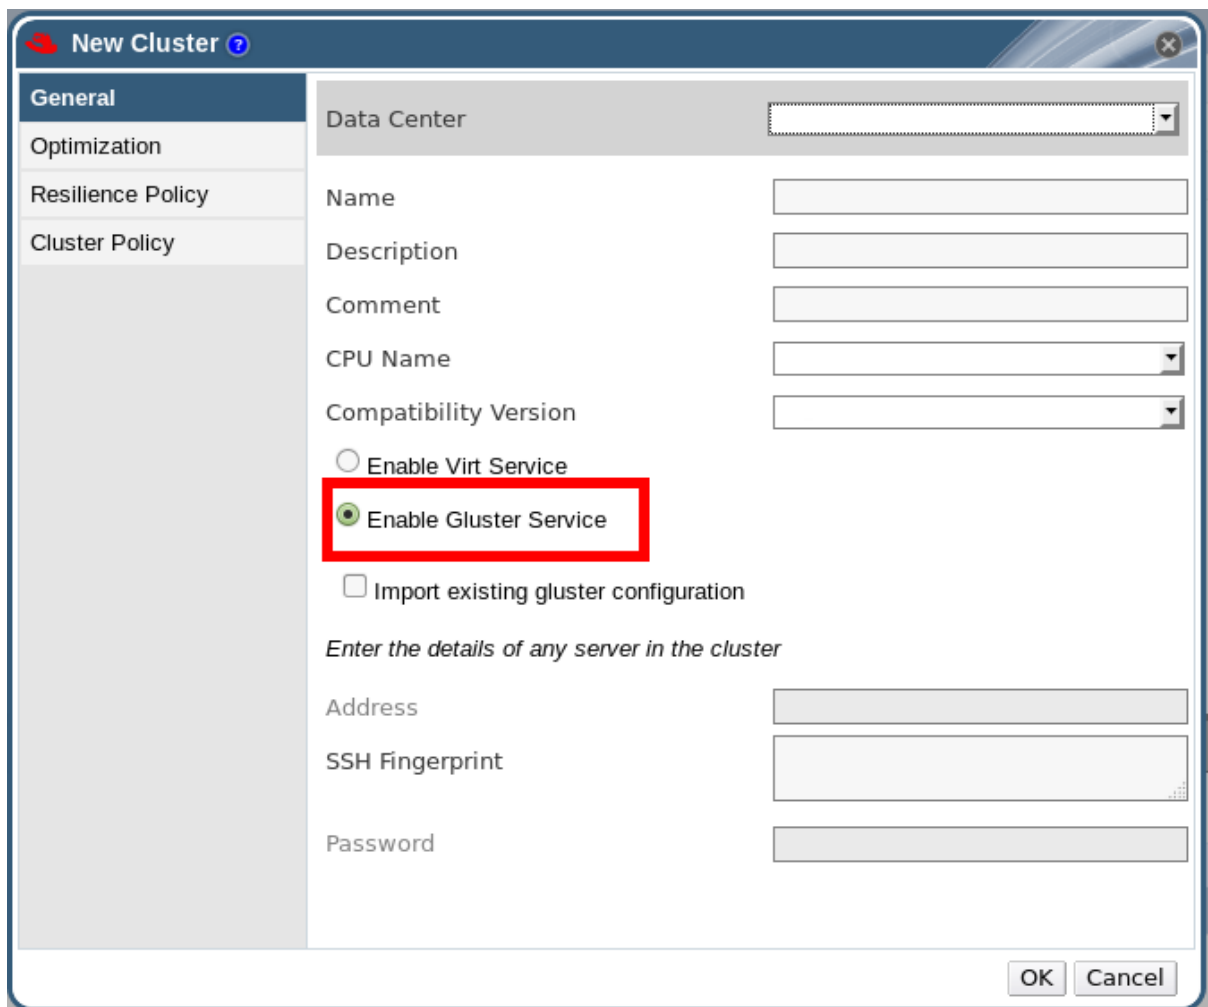


Figure 13.7. Selecting the "Enable Gluster Service" Radio Button

4. Click **OK**.

## Result

It is now possible to add Red Hat 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.

[Report a bug](#)

## 13.14. POPULATING THE ISO STORAGE DOMAIN

### Summary

An ISO storage domain is attached to a data center, ISO images must be uploaded to it. Red Hat Enterprise Virtualization provides an ISO uploader tool that ensures that the images are uploaded into the correct directory path, with the correct user permissions.

The creation of ISO images from physical media is not described in this document. It is assumed that you have access to the images required for your environment.

### Procedure 13.9. Populating the ISO Storage Domain

1. Copy the required ISO image to a temporary directory on the system running Red Hat Enterprise Virtualization Manager.
2. Log in to the system running Red Hat Enterprise Virtualization Manager as the **root** user.
3. Use the **engine-iso-uploader** command to upload the ISO image. This action will take some time, the amount of time varies depending on the size of the image being uploaded and available network bandwidth.

#### Example 13.1. ISO Uploader Usage

In this example the ISO image **RHEL6.iso** is uploaded to the ISO domain called **ISODomain** using NFS. The command will prompt for an administrative user name and password. The user name must be provided in the form *user name@domain*.

```
# engine-iso-uploader --iso-domain=ISODomain upload RHEL6.iso
```

### Result

The ISO image is uploaded and appears in the ISO storage domain specified. It is also available in the list of available boot media when creating virtual machines in the data center which the storage domain is attached to.

[Report a bug](#)

## 13.15. VIRTIO AND GUEST TOOL IMAGE FILES

The virtio-win ISO and Virtual Floppy Drive (VFD) images, which contain the VirtIO drivers for Windows virtual machines, and the rhev-tools-setup ISO, which contains the Red Hat Enterprise Virtualization Guest Tools for Windows virtual machines, are copied to an ISO storage domain upon installation and configuration of the domain.

These image files provide software that can be installed on virtual machines to improve performance and usability. The most recent virtio-win and rhev-tools-setup files can be accessed via the following symbolic links on the file system of the Red Hat Enterprise Virtualization Manager:

- `/usr/share/virtio-win/virtio-win.iso`
- `/usr/share/virtio-win/virtio-win_x86.vfd`
- `/usr/share/virtio-win/virtio-win_amd64.vfd`
- `/usr/share/rhev-guest-tools-iso/rhev-tools-setup.iso`

These image files must be manually uploaded to ISO storage domains that were not created locally by the installation process. Use the **engine-iso-uploader** command to upload these images to your ISO storage domain. Once uploaded, the image files can be attached to and used by virtual machines.

[Report a bug](#)

## 13.16. UPLOADING THE VIRTIO AND GUEST TOOL IMAGE FILES TO AN ISO STORAGE DOMAIN

The example below demonstrates the command to upload the **virtio-win.iso**, **virtio-win\_x86.vfd**, **virtio-win\_amd64.vfd**, and **rhev-tools-setup.iso** image files to the **ISODomain**.

### Example 13.2. Uploading the VirtIO and Guest Tool Image Files

```
# engine-iso-uploader --iso-domain=[ISODomain] upload /usr/share/virtio-win/virtio-win.iso
/usr/share/virtio-win/virtio-win_x86.vfd /usr/share/virtio-win/virtio-win_amd64.vfd /usr/share/rhev-
guest-tools-iso/rhev-tools-setup.iso
```

[Report a bug](#)

## CHAPTER 14. CONFIGURING LOGS

### 14.1. RED HAT ENTERPRISE VIRTUALIZATION MANAGER INSTALLATION LOG FILES

Table 14.1. Installation

Log File	Description
<code>/var/log/ovirt-engine/engine-cleanup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <b>engine-cleanup</b> command. This is the command used to reset a Red Hat Enterprise Virtualization Manager installation. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist.
<code>/var/log/ovirt-engine/engine-db-install-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <b>engine-setup</b> command detailing the creation and configuration of the <b>rhev</b> database.
<code>/var/log/ovirt-engine/rhev-dwh-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <b>rhev-dwh-setup</b> command. This is the command used to create the <b>ovirt_engine_history</b> database for reporting. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/ovirt-engine-reports-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the <b>rhev-reports-setup</b> command. This is the command used to install the Red Hat Enterprise Virtualization Manager Reports modules. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.
<code>/var/log/ovirt-engine/setup/ovirt-engine-setup-yyyymmddhhmmss.log</code>	Log from the <b>engine-setup</b> command. A log is generated each time the command is run. The date and time of the run is used in the filename to allow multiple logs to exist concurrently.

[Report a bug](#)

### 14.2. RED HAT ENTERPRISE VIRTUALIZATION MANAGER LOG FILES

Table 14.2. Service Activity

Log File	Description
<code>/var/log/ovirt-engine/engine.log</code>	Reflects all Red Hat Enterprise Virtualization Manager GUI crashes, Active Directory look-ups, Database issues, and other events.



Log File	Description
<b>/var/log/ovirt-engine/host-deploy</b>	Log files from hosts deployed from the Red Hat Enterprise Virtualization Manager.
<b>/var/lib/ovirt-engine/setup-history.txt</b>	Tracks the installation and upgrade of packages associated with the Red Hat Enterprise Virtualization Manager.

[Report a bug](#)

### 14.3. RED HAT ENTERPRISE VIRTUALIZATION HOST LOG FILES

Table 14.3.

Log File	Description
<b>/var/log/vdsm/libvirt.log</b>	Log file for <b>libvirt</b> .
<b>/var/log/vdsm/spm-lock.log</b>	Log file detailing the host's ability to obtain a lease on the Storage Pool Manager role. The log details when the host has acquired, released, renewed, or failed to renew the lease.
<b>/var/log/vdsm/vdsm.log</b>	Log file for VDSM, the Manager's agent on the virtualization host(s).
<b>/tmp/ovirt-host-deploy-@DATE@.log</b>	Host deployment log, copied to engine as <code>/var/log/ovirt-engine/host-deploy/ovirt-@DATE@-@HOST@-@CORRELATION_ID@.log</code> after the host has been successfully deployed.

[Report a bug](#)

### 14.4. SETTING UP A VIRTUALIZATION HOST LOGGING SERVER

#### Summary

Red Hat Enterprise Virtualization hosts generate and update log files, recording their actions and problems. Collecting these log files centrally simplifies debugging.

This procedure should be used on your centralized log server. You could use a separate logging server, or use this procedure to enable host logging on the Red Hat Enterprise Virtualization Manager.

#### Procedure 14.1. Setting up a Virtualization Host Logging Server

1. Configure SELinux to allow **rsyslog** traffic.

```
# semanage port -a -t syslogd_port_t -p udp 514
```

2. Edit `/etc/rsyslog.conf` and add below lines:

```
$template TmplAuth, "/var/log/%fromhost%/secure"  
$template TmplMsg, "/var/log/%fromhost%/messages"  
  
$RuleSet remote  
authpriv.* ?TmplAuth  
*.info,mail.none;authpriv.none,cron.none ?TmplMsg  
$RuleSet RSYSLOG_DefaultRuleset  
$InputUDPServerBindRuleset remote
```

Uncomment the following:

```
#$ModLoad imudp  
#$UDPServerRun 514
```

3. Restart the rsyslog service:

```
# service rsyslog restart
```

### Result

Your centralized log server is now configured to receive and store the **messages** and **secure** logs from your virtualization hosts.

[Report a bug](#)

## 14.5. THE LOGGING SCREEN

### Summary

The **Logging** screen allows you to configure logging-related options such as a daemon for automatically exporting log files generated by the hypervisor to a remote server.

### Procedure 14.2. Configuring Logging

1. In the **Logrotate Max Log Size** field, enter the maximum size in kilobytes that log files can reach before they are rotated by **logrotate**. The default value is **1024**.
2. Optionally, configure **rsyslog** to transmit log files to a remote **syslog** daemon:
  - a. Enter the remote **rsyslog** server address in the **Server Address** field.
  - b. Enter the remote **rsyslog** server port in the **Server Port** field. The default port is **514**.
3. Optionally, configure **netconsole** to transmit kernel messages to a remote destination:
  - a. Enter the **Server Address**.
  - b. Enter the **Server Port**. The default port is **6666**.
4. Select **<Save>** and press **Enter**.

### Result

You have configured logging for the hypervisor.

[Report a bug](#)

## PART V. ADVANCED SETUP

## CHAPTER 15. PROXIES

### 15.1. SPICE PROXY

#### 15.1.1. SPICE Proxy Overview

The SPICE Proxy is a tool used to connect SPICE Clients to guests when the SPICE Clients are outside the network that connects the hypervisors.

Setting up a SPICE Proxy consists of installing **Squid** on a machine and configuring **iptables** to allow proxy traffic through the firewall.

Turning a SPICE Proxy on consists of using **engine-config** on the Manager to set the key **SpiceProxyDefault** to a value consisting of the name and port of the proxy.

Turning a SPICE Proxy off consists of using **engine-config** on the Manager to remove the value that the key **SpiceProxyDefault** has been set to.

[Report a bug](#)

#### 15.1.2. SPICE Proxy Machine Setup

##### Summary

This procedure explains how to set up a machine as a SPICE Proxy. A SPICE Proxy makes it possible to connect to the Red Hat Enterprise Virtualization network from outside the network. We use **Squid** in this procedure to provide proxy services.

##### Procedure 15.1. Installing Squid on Red Hat Enterprise Linux

1. Install **Squid** on the Proxy machine:

```
# yum install squid
```

2. Open **/etc/squid/squid.conf**. Change:

```
http_access deny CONNECT !SSL_ports
```

to:

```
http_access deny CONNECT !Safe_ports
```

3. Restart the proxy:

```
# service squid restart
```

4. Open the default squid port:

```
# iptables -A INPUT -p tcp --dport 3128 -j ACCEPT
```

5. Make this iptables rule persistent:

```
# service iptables save
```

## Result

You have now set up a machine as a SPICE proxy. Before connecting to the Red Hat Enterprise Virtualization network from outside the network, activate the SPICE proxy.

[Report a bug](#)

### 15.1.3. Turning on SPICE Proxy

#### Summary

This procedure explains how to activate (or turn on) the SPICE proxy.

#### Procedure 15.2. Activating SPICE Proxy

1. On the Manager, use the `engine-config` tool to set a proxy:


```
# engine-config -s SpiceProxyDefault=someProxy
```

2. Restart the `ovirt-engine` service:

```
# service ovirt-engine restart
```

The proxy must have this form:

```
protocol://[host]:[port]
```



#### NOTE

Only the `http` protocol is supported by SPICE clients. If `https` is specified, the client will ignore the proxy setting and attempt a direct connection to the hypervisor.

## Result

SPICE Proxy is now activated (turned on). It is now possible to connect to the Red Hat Enterprise Virtualization network through the SPICE proxy.

[Report a bug](#)

### 15.1.4. Turning Off a SPICE Proxy

#### Summary

This procedure explains how to turn off (deactivate) a SPICE proxy.

#### Procedure 15.3. Turning Off a SPICE Proxy

1. Log in to the Manager:

```
$ ssh root@[IP of Manager]
```

2. Run the following command to clear the SPICE proxy:

```
# engine-config -s SpiceProxyDefault=""
```

3. Restart the Manager:

```
# service ovirt-engine restart
```

### Result

SPICE proxy is now deactivated (turned off). It is no longer possible to connect to the Red Hat Enterprise Virtualization network through the SPICE proxy.

[Report a bug](#)

## 15.2. SQUID PROXY

### 15.2.1. Installing and Configuring a Squid Proxy

#### Summary

This section explains how to install and configure a Squid Proxy to the User Portal.

#### Procedure 15.4. Configuring a Squid Proxy

1. **Obtaining a Keypair**

Obtain a keypair and certificate for the HTTPS port of the Squid proxy server.

You can obtain this keypair the same way that you would obtain a keypair for another SSL/TLS service. The keypair is in the form of two PEM files which contain the private key and the signed certificate. In this document we assume that they are named **proxy.key** and **proxy.cer**.

The keypair and certificate can also be generated using the certificate authority of the oVirt engine. If you already have the private key and certificate for the proxy and do not want to generate it with the oVirt engine certificate authority, skip to the next step.

2. **Generating a Keypair**

Decide on a host name for the proxy. In this procedure, the proxy is called **proxy.example.com**.

Decide on the rest of the distinguished name of the certificate for the proxy. The important part here is the "common name", which contains the host name of the proxy. Users' browsers use the common name to validate the connection. It is good practice to use the same country and same organization name used by the oVirt engine itself. Find this information by logging in to the oVirt engine machine and running the following command:

```
[root@engine ~]# openssl x509 -in /etc/pki/ovirt-engine/ca.pem -noout -subject
```

This command will output something like this:

```
subject= /C=US/O=Example Inc./CN=engine.example.com.81108
```

The relevant part here is **/C=us/O=Example Inc..** Use this to build the complete distinguished name for the certificate for the proxy:

```
/C=US/O=Example Inc./CN=proxy.example.com
```

Log in to the proxy machine and generate a certificate signing request:

```
[root@proxy ~]# openssl req -newkey rsa:2048 -subj '/C=US/O=Example
Inc./CN=proxy.example.com' -nodes -keyout proxy.key -out proxy.req
```



#### NOTE

The quotes around the distinguished name for the certificate are very important. Do not leave them out.

The command will generate the key pair. It is very important that the private key is not encrypted (that is the effect of the `-nodes` option) because otherwise you would need to type the password to start the proxy server.

The output of the command looks like this:

```
Generating a 2048 bit RSA private key
.....+++
.....+++
writing new private key to 'proxy.key'
-----
```

The command will generate two files: **proxy.key** and **proxy.req**. **proxy.key** is the private key. Keep this file safe. **proxy.req** is the certificate signing request. **proxy.req** does not require any special protection.

To generate the signed certificate, copy the **private.csr** file to the oVirt engine machine, using the **scp** command:

```
[root@proxy ~]# scp proxy.req engine.example.com:/etc/pki/ovirt-engine/requests/
```

Log in to the oVirt engine machine and run the following command to sign the certificate:

```
[root@engine ~]# /usr/share/ovirt-engine/bin/pki-enroll-request.sh --name=proxy --
days=3650 --subject='/C=US/O=Example Inc./CN=proxy.example.com'
```

This will sign the certificate and make it valid for 10 years (3650 days). Set the certificate to expire earlier, if you prefer.

The output of the command looks like this:

```
Using configuration from openssl.conf
Check that the request matches the signature
Signature ok
The Subject's Distinguished Name is as follows
countryName      :PRINTABLE:'US'
organizationName :PRINTABLE:'Example Inc.'
commonName       :PRINTABLE:'proxy.example.com'
Certificate is to be certified until Jul 10 10:05:24 2023 GMT (3650
days)
```



```
Write out database with 1 new entries
Data Base Updated
```

The generated certificate file is available in the directory `/etc/pki/ovirt-engine/certs` and should be named **proxy.cer**. Copy this file to the proxy machine:

```
[root@proxy ~]# scp engine.example.com:/etc/pki/ovirt-engine/certs/proxy.cer .
```

Make sure that both the **proxy.key** and **proxy.cer** files are present on the proxy machine:

```
[root@proxy ~]# ls -l proxy.key proxy.cer
```

The output of this command will look like this:

```
-rw-r--r--. 1 root root 4902 Jul 12 12:11 proxy.cer
-rw-r--r--. 1 root root 1834 Jul 12 11:58 proxy.key
```

You are now ready to install and configure the proxy server.

### 3. Install the Squid proxy server package

Install this system as follows:

```
[root@proxy ~]# yum -y install squid
```

### 4. Configure the Squid proxy server

Move the private key and signed certificate to a place where the proxy can access them, for example to the `/etc/squid` directory:

```
[root@proxy ~]# cp proxy.key proxy.cer /etc/squid/.
```

Set permissions so that the "squid" user can read these files:

```
[root@proxy ~]# chgrp squid /etc/squid/proxy.*
[root@proxy ~]# chmod 640 /etc/squid/proxy.*
```

The Squid proxy will connect to the oVirt engine web server using the SSL protocol, and must verify the certificate used by the engine. Copy the certificate of the CA that signed the certificate of the oVirt engine web server to a place where the proxy can access it, for example `/etc/squid`. The default CA certificate is located in the `/etc/pki/ovirt-engine/ca.pem` file in the oVirt engine machine. Copy it with the following command:

```
[root@proxy ~]# scp engine.example.com:/etc/pki/ovirt-engine/ca.pem /etc/squid/.
```

Ensure the **squid** user can read that file:

```
[root@proxy ~]# chgrp squid /etc/squid/ca.pem
[root@proxy ~]# chmod 640 /etc/squid/ca.pem
```

If SELinux is in enforcing mode, change the context of port 443 using the **semanage** tool. This permits Squid to use port 443.

```
[root@proxy ~]# yum install -y polycoreutils-python
[root@proxy ~]# semanage port -m -p tcp -t http_cache_port_t 443
```

Replace the existing squid configuration file with the following:

```
https_port 443 key=/etc/squid/proxy.key cert=/etc/squid/proxy.cer ssl-bump
defaultsite=engine.example.com
cache_peer engine.example.com parent 443 0 no-query originserver ssl
sslcafile=/etc/squid/ca.pem name=engine
cache_peer_access engine allow all
ssl_bump allow all
http_access allow all
```

## 5. Restart the Squid Proxy Server

Run the following command in the proxy machine:

```
[root@proxy ~]# service squid restart
```

## 6. Configure the websockets proxy



### NOTE

This step is optional. Do this step only to use the noVNC console or the SPICE HTML 5 console.

To use the noVNC or SPICE HTML 5 consoles to connect to the console of virtual machines, the websocket proxy server must be configured on the machine on which the engine is installed. If you selected to configure the websocket proxy server when prompted during installing or upgrading the engine with the **engine-setup** command, the websocket proxy server will already be configured. If you did not select to configure the websocket proxy server at this time, you can configure it later by running the **engine-setup** command with the following option:


```
engine-setup --otopi-environment="OVESETUP_CONFIG/websocketProxyConfig=bool:True"
```

You must also ensure the **ovirt-websocket-proxy** service is started and will start automatically on boot:

```
[root@engine ~]# service ovirt-websocket-proxy status
[root@engine ~]# chkconfig ovirt-websocket-proxy on
```

Both the noVNC and the SPICE HTML 5 consoles use the websocket protocol to connect to the virtual machines, but squid proxy server does not support the websockets protocol, so this communication cannot be proxied with Squid. Tell the system to connect directly to the websockets proxy running in the machine where the engine is running. To do this, update the **WebSocketProxy** configuration parameter using the "engine-config" tool:

```
[root@engine ~]# engine-config \
-s WebSocketProxy=engine.example.com:6100
[root@engine ~]# service ovirt-engine restart
```



## IMPORTANT

If you skip this step the clients will assume that the websockets proxy is running in the proxy machine, and thus will fail to connect.

### 7. Connect to the user portal using the complete URL

Connect to the User Portal using the complete URL, for instance:

```
https://proxy.example.com/UserPortal/org.ovirt.engine.ui.userportal.UserPortal/UserPortal.html
```



## NOTE

Shorter URLs, for example **https://proxy.example.com/UserPortal**, will not work. These shorter URLs are redirected to the long URL by the application server, using the 302 response code and the Location header. The version of **Squid** in Red Hat Enterprise Linux and Fedora ( **Squid** version 3.1) does not support rewriting these headers.

## Summary

You have installed and configured a Squid proxy to the User Portal.

[Report a bug](#)

## APPENDIX A. REVISION HISTORY

Revision 3.4-36	Fri 03 Jul 2015	Red Hat Enterprise Virtualization Documentation Team
<a href="#">BZ#1232136</a> - New section 'Backing up and Restoring a Self-Hosted Environment' added with six new topics.		
Revision 3.4-35	Wed 13 May 2015	Red Hat Enterprise Virtualization Documentation Team
<a href="#">BZ#1160742</a> - Updated the output for configuring Data Warehouse and Reports, and added a note about moving self-hosted engine environments to maintenance mode.		
Revision 3.4-34	Tue 14 Apr 2015	Red Hat Enterprise Virtualization Documentation Team
<a href="#">BZ#1195104</a> - Updated information for the three maintenance modes of the hosted engine.		
Revision 3.4-33	Fri 20 Mar 2015	Red Hat Enterprise Virtualization Documentation Team
<a href="#">BZ#1203487</a> - Removed references to now-defunct Hypervisor Deployment Guide.		
Revision 3.4-32	Thu 05 Mar 2015	Red Hat Enterprise Virtualization Documentation Team
<a href="#">BZ#1122912</a> - Added documentation for the Diagnostics screen of the hypervisor configuration menu.		
<a href="#">BZ#1122915</a> - Added documentation for the Performance screen of the hypervisor configuration menu.		
<a href="#">BZ#1122919</a> - Added documentation for the Plugins screen of the hypervisor configuration menu.		
Revision 3.4-31	Thu 05 Feb 2015	Andrew Dahms
Updated the supported version of Red Hat Enterprise Linux.		
Revision 3.4-30	Thurs 11 Dec 2014	Tahlia Richardson
<a href="#">BZ#1172299</a> - Updated the command for saving iptables rules persistently.		
Revision 3.4-29	Tues 28 Oct 2014	Tahlia Richardson
<a href="#">BZ#1157713</a> - Removed the default storage type question from RHEV-M setup.		
Revision 3.4-28	Tue 28 Oct 2014	Julie Wu
<a href="#">BZ#1154537</a> - Added an important note on upgrading to the latest minor version before upgrading to the next major version.		
Revision 3.4-27	Tue 07 Oct 2014	Julie Wu
<a href="#">BZ#1145040</a> - Added a note referencing the RHEL Security Guide.		
Revision 3.4-26	Thu 28 Aug 2014	Andrew Dahms
<a href="#">BZ#1083382</a> - Added a section outlining how to modify the Red Hat Enterprise Virtualization Hypervisor ISO file.		
<a href="#">BZ#853119</a> - Added a description of how to modify user and group IDs in the Red Hat Enterprise Virtualization Hypervisor ISO file.		
Revision 3.4-25	Mon 25 Aug 2014	Julie Wu
<a href="#">BZ#1123739</a> - Updated the kbase article link for offline installation.		
Revision 3.4-24	Tue 15 Jul 2014	Andrew Burden
<a href="#">BZ#1104114</a> - 'Installing the Self-Hosted Engine' now clearly lists the channels required to install the ovirt-self-hosted package.		
Revision 3.4-23	Fri 13 Jun 2014	Zac Dover
rhev-m-doc rebuild		
Revision 3.4-22	Wed 11 Jun 2014	Andrew Burden
Brewhing for 3.4 GA		

<b>Revision 3.4-21</b>	<b>Tue 10 Jun 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1107996</a> - Updated the port numbers used by libvirt.		
<a href="#">BZ#1094069</a> - Updated the procedure for upgrading to Red Hat Enterprise Virtualization 3.4.		
<a href="#">BZ#1075942</a> - Updated the options for the engine-cleanup command.		
<b>Revision 3.4-20</b>	<b>Wed 30 Apr 2014</b>	<b>Zac Dover</b>
Final build.		
<b>Revision 3.4-19</b>	<b>Tue 29 Apr 2014</b>	<b>Andrew Burden</b>
<a href="#">BZ#1092075</a> - Updated 'Upgrading the Self-Hosted Engine'		
<b>Revision 3.4-18</b>	<b>Mon 28 Apr 2014</b>	<b>Andrew Burden</b>
Added new topic 'Upgrading Additional Hosts in a Self-Hosted Environment'		
<b>Revision 3.4-17</b>	<b>Sun 27 Apr 2014</b>	<b>Andrew Burden</b>
<a href="#">BZ#1091576</a> - Added 'Upgrading the Self-Hosted Engine' topic		
<b>Revision 3.4-16</b>	<b>Wed 23 Apr 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1090715</a> - Updated the procedure for updating the guest tools.		
<a href="#">BZ#1090678</a> - Updated the procedures for preparing hypervisor installation media.		
<a href="#">BZ#1075418</a> - Updated the version of JBoss Enterprise Application Platform required for Red Hat Enterprise Virtualization.		
<b>Revision 3.4-16</b>	<b>Wed 23 Apr 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1090514</a> - Updated the host compatibility table to include Red Hat Enterprise Virtualization 3.4.		
<a href="#">BZ#1090480</a> - Updated the description regarding the limitations of using logical networks offered by external providers.		
<a href="#">BZ#1089856</a> - Updated the procedure for installing the Red Hat Enterprise Virtualization Hypervisor.		
<a href="#">BZ#1089871</a> - Updated the procedures for configuring Red Hat Enterprise Virtualization Hypervisors.		
<a href="#">BZ#1089762</a> - Updated the description of display ports that virtual machines can use.		
<b>Revision 3.4-15</b>	<b>Tue 22 Apr 2014</b>	<b>Lucy Bopf</b>
<a href="#">BZ#1076928</a> - Added drac7 as an option for power management device.		
<a href="#">BZ#1076926</a> - Added hplblade as an option for power management device.		
<b>Revision 3.4-14</b>	<b>Thu 17 Apr 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1089762</a> - Updated the description of ports that must be enabled on hosts.		
<a href="#">BZ#1088666</a> - Removed all references to beta channels.		
<a href="#">BZ#1087691</a> - Updated the channels required to install packages for Red Hat Enterprise Virtualization 3.4.		
<b>Revision 3.4-13</b>	<b>Thu 17 Apr 2014</b>	<b>Lucy Bopf</b>
<a href="#">BZ#1075519</a> - Added note that Multi-Host Network Configuration is now active when editing logical networks.		
<b>Revision 3.4-12</b>	<b>Thu 10 Apr 2014</b>	<b>Lucy Bopf</b>
<a href="#">BZ#1076274</a> - Updated Storage Domain information to reflect the fact that storage domains of multiple types can be added to the same data center.		
<a href="#">BZ#1075909</a> - Updated procedure for upgrading data centers to include confirmation window and warning.		
<a href="#">BZ#1075532</a> - Updated Cluster Policy information to include new power off capacity.		
<a href="#">BZ#1025433</a> - Added a note that detailed guides can be found in the JasperReports subfolder.		
<a href="#">BZ#1075253</a> - Updated procedure for creating a new cluster to include new functions.		
<a href="#">BZ#1075538</a> - Updated procedure for creating a new cluster to include new Enable HA Reservation option.		
<b>Revision 3.4-11</b>	<b>Fri 04 Apr 2014</b>	<b>Andrew Dahms</b>
<a href="#">BZ#1088086</a> - Updated the description of features requiring a compatibility upgrade to version 3.4.		
<a href="#">BZ#1087646</a> - Updated the procedure for installing the Red Hat Enterprise Virtualization Manager.		
<a href="#">BZ#1083848</a> - Updated the procedure for upgrading Red Hat Enterprise Virtualization.		
<a href="#">BZ#1083768</a> - Changed instances of 'yum config-manager' to 'yum-config-manager'.		
<a href="#">BZ#1059543</a> - Corrected the arguments to the hosted-engine command for setting the maintenance mode.		

