



Red Hat Enterprise Virtualization 3.5 Installation Guide

Installing Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization Documentation Team

Red Hat Enterprise Virtualization 3.5 Installation Guide

Installing Red Hat Enterprise Virtualization

Red Hat Enterprise Virtualization Documentation Team
Red Hat Customer Content Services
rhev-docs@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](#). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

A comprehensive guide to installing Red Hat Enterprise Virtualization.

Table of Contents

Part I. Introduction	4
Chapter 1. Introduction	5
1.1. Workflow Progress - System Requirements	5
1.2. Red Hat Enterprise Virtualization Manager Requirements	5
1.3. Hypervisor Requirements	7
1.4. User Authentication	11
1.5. Firewalls	13
1.6. System Accounts	18
Part II. Installing Red Hat Enterprise Virtualization	20
Chapter 2. Installing Red Hat Enterprise Virtualization	21
2.1. Workflow Progress - Installing Red Hat Enterprise Virtualization Manager	21
2.2. Overview of Installing the Red Hat Enterprise Virtualization Manager	21
2.3. Subscribing to the Required Entitlements	24
2.4. Installing the Red Hat Enterprise Virtualization Manager	25
2.5. SPICE Client	36
Chapter 3. The Self-Hosted Engine	38
3.1. About the Self-Hosted Engine	38
3.2. Subscribing to the Required Entitlements	38
3.3. Installing the Self-Hosted Engine	39
3.4. Configuring the Self-Hosted Engine	40
3.5. Installing Additional Hosts to a Self-Hosted Environment	47
3.6. Maintaining the Self-Hosted Engine	49
3.7. Upgrading the Self-Hosted Engine	49
3.8. Upgrading Additional Hosts in a Self-Hosted Environment	51
3.9. Removing a Host from a Self-Hosted Engine Environment	52
3.10. Backing up and Restoring a Self-Hosted Engine Environment	53
3.11. Migrating to a Self-Hosted Environment	75
3.12. Migrating the Self-Hosted Engine Database to a Remote Server Database	83
3.13. Migrating the Data Warehouse Database to a Remote Server Database	86
3.14. Migrating the Reports Database to a Remote Server Database	89
Chapter 4. Data Warehouse and Reports	93
4.1. Workflow Progress - Data Collection Setup and Reports Installation	93
4.2. Overview of Configuring Data Warehouse and Reports	93
4.3. Data Warehouse and Reports Configuration Notes	94
4.4. Data Warehouse and Reports Installation Options	94
4.5. Migrating Data Warehouse and Reports to Separate Machines	116
Chapter 5. Updating the Red Hat Enterprise Virtualization Environment	126
5.1. Updates between Minor Releases	126
5.2. Upgrading to Red Hat Enterprise Virtualization 3.5	130
5.3. Upgrading to Red Hat Enterprise Virtualization 3.4	133
5.4. Upgrading to Red Hat Enterprise Virtualization 3.3	136
5.5. Upgrading to Red Hat Enterprise Virtualization Manager 3.2	140
5.6. Upgrading to Red Hat Enterprise Virtualization Manager 3.1	142
5.7. Post-Upgrade Tasks	145
Part III. Installing Hosts	147
Chapter 6. Introduction to Hosts	148

6.1. Workflow Progress - Installing Virtualization Hosts	148
6.2. Introduction to Virtualization Hosts	148
Chapter 7. Red Hat Enterprise Virtualization Hypervisor Hosts	150
7.1. Red Hat Enterprise Virtualization Hypervisor Installation Overview	150
7.2. Obtaining the Red Hat Enterprise Virtualization Hypervisor Disk Image	150
7.3. Preparing Installation Media	151
7.4. Installing the Red Hat Enterprise Virtualization Hypervisor	155
7.5. Automated Installation	160
7.6. Configuration	173
7.7. Adding Hypervisors to Red Hat Enterprise Virtualization Manager	190
7.8. Modifying the Red Hat Enterprise Virtualization Hypervisor ISO	191
Chapter 8. Red Hat Enterprise Linux Hosts	199
8.1. Red Hat Enterprise Linux Hosts	199
8.2. Host Compatibility Matrix	199
8.3. Installing Red Hat Enterprise Linux	199
8.4. Subscribing to the Required Entitlements	202
8.5. Configuring the Virtualization Host Firewall	203
8.6. Configuring Virtualization Host sudo	204
8.7. Configuring Virtualization Host SSH	205
8.8. Adding a Red Hat Enterprise Linux Host	206
8.9. Explanation of Settings and Controls in the New Host and Edit Host Windows	207
Part IV. Basic Setup	212
Chapter 9. Configuring Data Centers	213
9.1. Workflow Progress - Planning Your Data Center	213
9.2. Planning Your Data Center	213
9.3. Data Centers in Red Hat Enterprise Virtualization	214
9.4. Creating a New Data Center	214
9.5. Changing the Data Center Compatibility Version	215
Chapter 10. Configuring Clusters	217
10.1. Clusters in Red Hat Enterprise Virtualization	217
10.2. Creating a New Cluster	217
10.3. Changing the Cluster Compatibility Version	218
Chapter 11. Configuring Networking	220
11.1. Workflow Progress - Network Setup	220
11.2. Networking in Red Hat Enterprise Virtualization	220
11.3. Creating Logical Networks	222
11.4. Editing Logical Networks	223
11.5. External Provider Networks	231
11.6. Bonding	234
11.7. Removing Logical Networks	237
Chapter 12. Configuring Storage	239
12.1. Workflow Progress - Storage Setup	239
12.2. Introduction to Storage in Red Hat Enterprise Virtualization	239
12.3. Preparing NFS Storage	240
12.4. Attaching NFS Storage	242
12.5. Changing the Permissions for the Local ISO Domain	243
12.6. Attaching the Local ISO Domain to a Data Center	243
12.7. Adding iSCSI Storage	244

12.8. Adding FCP Storage	246
12.9. Preparing Local Storage	248
12.10. Adding Local Storage	248
12.11. POSIX Compliant File System Storage in Red Hat Enterprise Virtualization	250
12.12. Attaching POSIX Compliant File System Storage	250
12.13. Enabling Gluster Processes on Red Hat Gluster Storage Nodes	252
12.14. Populating the ISO Storage Domain	252
12.15. VirtIO and Guest Tool Image Files	253
12.16. Uploading the VirtIO and Guest Tool Image Files to an ISO Storage Domain	253
Chapter 13. Configuring Logs	255
13.1. Red Hat Enterprise Virtualization Manager Installation Log Files	255
13.2. Red Hat Enterprise Virtualization Manager Log Files	255
13.3. Red Hat Enterprise Virtualization Host Log Files	255
13.4. Setting Up a Virtualization Host Logging Server	256
13.5. The Logging Screen	257
Part V. Advanced Setup	258
Chapter 14. Proxies	259
14.1. SPICE Proxy	259
14.2. Squid Proxy	261
Appendix A. Red Hat Enterprise Virtualization Installation Options	265
A.1. Configuring a Local Repository for Offline Red Hat Enterprise Virtualization Manager Installation	265
A.2. Updating the Local Repository for an Offline Red Hat Enterprise Virtualization Manager Installation	267
Appendix B. Revision History	268

Part I. Introduction

Chapter 1. Introduction

1.1. Workflow Progress - System Requirements



1.2. Red Hat Enterprise Virtualization Manager Requirements

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

- ✧ One machine to act as the Red Hat Enterprise Virtualization Manager.
- ✧ 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 act 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, or FC, or be locally attached to virtualization hosts. The use of other POSIX compliant filesystems is also supported.

1.2.1. 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. These recommendations are a guide only.

Minimum

- » A dual core CPU.
- » 4 GB of available system RAM if Data Warehouse is not installed and if memory is not being consumed by existing processes.
- » 25 GB of locally accessible, writable, 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, writable, 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, see <https://hardware.redhat.com/>.

1.2.2. Operating System Requirements

The Red Hat Enterprise Virtualization Manager must run on Red Hat Enterprise Linux Server 6.5, 6.6, or 6.7. 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.

1.2.3. Browser and Client Requirements

The following browser versions and operating systems have supported SPICE clients and are

optimal for displaying the application graphics of the Administration Portal and the User Portal:

Operating System Family	Browser	Portal Access	Supported SPICE Client?
Red Hat Enterprise Linux	Mozilla Firefox 38	Administration Portal and User Portal	Yes
Windows	Internet Explorer 9 or later	Administration Portal	Yes
	Internet Explorer 8 or later	User Portal	Yes

1.2.4. Software Repositories

To install Red Hat Enterprise Virtualization Manager, you must have the following entitlements before proceeding with installation.

Table 1.1. Required Pools for Red Hat Enterprise Virtualization Manager

Subscription pool	Repository name	Repository label	Details
Red Hat Enterprise Linux Server	Red Hat Enterprise Linux Server	rhel-6-server-rpms	Provides the Red Hat Enterprise Linux 6 Server.
Red Hat Enterprise Linux Server	RHEL Server Supplementary	rhel-6-server-supplementary-rpms	Provides the <i>virtio-win</i> package, which provides the Windows VirtIO drivers for use in virtual machines.
Red Hat Enterprise Virtualization	Red Hat Enterprise Virtualization	rhel-6-server-rhev-3.5-rpms	Provides the Red Hat Enterprise Virtualization Manager.
Red Hat Enterprise Virtualization	Red Hat JBoss Enterprise Application Platform	jb-eap-6-for-rhel-6-server-rpms	Provides the supported release of Red Hat JBoss Enterprise Application Platform on which the Manager runs.

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.

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
- » IBM POWER8

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.



Note

You must enable Virtualization in the BIOS. Cold boot the host after this change to ensure that the change is applied.

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 the following factors:

- » Guest operating system requirements.
- » Guest application requirements.

- » Memory activity and usage of guests.

The fact that KVM is able to overcommit 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.

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.

For Red Hat Enterprise Virtualization Hypervisor requirements, see the following table for the minimum supported internal storage for each version of the Hypervisor:

Table 1.2. Red Hat Enterprise Virtualization Hypervisor Minimum Storage Requirements

Version	Root Partition	Configuration Partition	Logging Partition	Data Partition	Swap Partition	Minimum Total
Red Hat Enterprise Virtualization Hypervisor 6	512 MB	8 MB	2048 MB	512 MB	8 MB	3.5 GB
Red Hat Enterprise Virtualization Hypervisor 7	9 GB	8 MB	2048 MB	512 MB	8 MB	12 GB

The logging partition requires a minimum of 2048 MB storage. However, it is recommended to allocate more storage to the logging partition if resources permit.

The data partition requires a minimum of 512 MB storage. The recommended size is at least 1.5 times as large as the RAM on the host system plus an additional 512 MB. 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.

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 one dedicated to support network intensive activities such as virtual machine migration. The performance of such operations are limited by the bandwidth available.

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.

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. For directory servers implemented prior to Red Hat Enterprise Virtualization 3.5, use the Domain Management Tool with the **engine-manage-domains** command to manage your domains. See the [The Domain Management Tool](#) section of the *Red Hat Enterprise Virtualization Administration Guide* for more information. With Red Hat Enterprise Virtualization 3.5, use the new generic LDAP provider implementation. See [Configuring a Generic LDAP Provider](#) section of the *Red Hat Enterprise Virtualization Administration Guide* for more information.

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 3.5 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 lookup 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
- ✦ Red Hat Directory Server (RHDS) - http://docs.redhat.com/docs/en-US/Red_Hat_Directory_Server/index.html
- ✦ OpenLDAP - <http://www.openldap.org/doc/>



Important

A user with permissions to browse all users and groups 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. The Red Hat Directory Server (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.

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 can 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 1.3. Red Hat Enterprise Virtualization Manager Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
-	ICMP	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Red Hat Enterprise Virtualization Manager	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	System(s) used for maintenance of the Manager including backend configuration, and software upgrades.	Red Hat Enterprise Virtualization Manager	Secure Shell (SSH) access. Optional.
80, 443	TCP	Administration Portal clients User Portal clients Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s) REST API clients	Red Hat Enterprise Virtualization Manager	Provides HTTP and HTTPS access to the Manager.
6100	TCP	Administration Portal clients User Portal clients	Red Hat Enterprise Virtualization Manager	Provides websocket proxy access for web-based console clients (noVNC and spice-html5) when the websocket proxy is running on the Manager. If the websocket proxy is running on a different host, however, this port is not used.

Port(s)	Protocol	Source	Destination	Purpose
7410	UDP	Red Hat Enterprise Virtualization Hypervisor(s)	Red Hat Enterprise Virtualization Manager	Must be open for the Manager to receive Kdump notifications.
		Red Hat Enterprise Linux host(s)		



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_TCPPORT="port"**
- ✧ UDP port specified with **LOCKD_UDPPORT="port"**

The **MOUNTD_PORT**, **STATD_PORT**, **LOCKD_TCPPORT**, and **LOCKD_UDPPORT** ports are configured in the **/etc/sysconfig/nfs** file.

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 1.4. Virtualization Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
22	TCP	Red Hat Enterprise Virtualization Manager	Red Hat Enterprise Virtualization Hypervisor(s)	Secure Shell (SSH) access. Optional.
			Red Hat Enterprise Linux host(s)	

Port(s)	Protocol	Source	Destination	Purpose
161	UDP	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Red Hat Enterprise Virtualization Manager	Simple network management protocol (SNMP). Only required if you want Simple Network Management Protocol traps sent from the hypervisor to one or more external SNMP managers. Optional.
5900 - 6923	TCP	Administration Portal clients User Portal clients	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Remote guest console access via VNC and SPICE. These ports must be open to facilitate client access to virtual machines.
5989	TCP, UDP	Common Information Model Object Manager (CIMOM)	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Used by Common Information Model Object Managers (CIMOM) to monitor virtual machines running on the hypervisor. Only required if you want to use a CIMOM to monitor the virtual machines in your virtualization environment. Optional.
16514	TCP	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Red Hat Enterprise Virtualization Hypervisor(s) Red Hat Enterprise Linux host(s)	Virtual machine migration using libvirt .

Port(s)	Protocol	Source	Destination	Purpose
49152 - 49216	TCP	Red Hat Enterprise Virtualization Hypervisor(s)	Red Hat Enterprise Virtualization Hypervisor(s)	Virtual machine migration and fencing using VDSM. These ports must be open facilitate both automated and manually initiated migration of virtual machines.
		Red Hat Enterprise Linux host(s)	Red Hat Enterprise Linux host(s)	
54321	TCP	Red Hat Enterprise Virtualization Manager	Red Hat Enterprise Virtualization Hypervisor(s)	VDSM communications with the Manager and other virtualization hosts.
		Red Hat Enterprise Virtualization Hypervisor(s)	Red Hat Enterprise Linux host(s)	
		Red Hat Enterprise Linux host(s)		

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 1.5. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
88, 464	TCP, UDP	Red Hat Enterprise Virtualization Manager	Directory server	Kerberos authentication.
389, 636	TCP	Red Hat Enterprise Virtualization Manager	Directory server	Lightweight Directory Access Protocol (LDAP) and LDAP over SSL.

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 1.6. Host Firewall Requirements

Port(s)	Protocol	Source	Destination	Purpose
5432	TCP, UDP	Red Hat Enterprise Virtualization Manager	PostgreSQL database server	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.

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 **vdsm** 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.

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 **vdsm** user.
- ✦ The **ovirt** group (GID **108**). Group members include:
 - The **ovirt** user.

1.6.3. Virtualization Host User Accounts

When the *vdsm* 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 *vdsm* and *qemu-kvm-rhev* package installation. The default user identifier (UID) for each entry is also listed:

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

In addition Red Hat Enterprise Virtualization Hypervisor hosts define an **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 **vdsm** 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 **vdsm** 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 *vdsm* and *qemu-kvm-rhev* packages.

1.6.4. Virtualization Host Groups

When the *vdsm* 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 *vdsm* 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 **vdsm** user.
 - The **sanlock** user.



Important

The user identifiers (UIDs) and group identifiers (GIDs) allocated may vary between systems. The **vdsm** 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 *vdsm* and *qemu-kvm-rhev* packages.

Part II. Installing Red Hat Enterprise Virtualization

Chapter 2. Installing Red Hat Enterprise Virtualization

2.1. Workflow Progress - Installing Red Hat Enterprise Virtualization Manager



2.2. Overview of 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 firewall rules, if any, present on the system. The default option is to allow the Manager's setup script to configure the firewall automatically; this overwrites any existing settings. To integrate the existing settings with the firewall rules required by the Manager, you must configure the firewall manually. If you choose to manually configure the firewall, the setup script provides a custom list of ports that need to be opened, based on the options selected during setup.
2. The fully qualified domain name (FQDN) of the system on which the Manager is to be installed. The default value is the system's current host name.
3. The password you use to secure the Red Hat Enterprise Virtualization administration account.
4. The location of the database server to be used as the Manager database. You can use the setup script to install and configure a local database server; this is the default setting. Alternatively, use an existing remote database server. This database must be created before the Manager is configured. To use a remote database server you must know:
 - ✦ The host name of the system on which the remote database server exists. The default host is **localhost**.
 - ✦ The port on which the remote database server is listening. The default port is 5432.
 - ✦ That the **uuid-oss** extension had been loaded by the remote database server.

You must also know the name of the database, and the user name and password of a user that has permissions on the remote database server. The default name for both the database and the user is **engine**.
5. The organization name to use when creating the Manager's security certificates. The default value is an automatically-detected domain-based name.
6. The following details about the local ISO domain, if the Manager is being configured to provide one:
 - ✦ The path for the ISO domain. The default path is **/var/lib/exports/iso**.
 - ✦ The networks or specific hosts that require access to the ISO domain. By default, the access control list (ACL) for the ISO domain provides read and write access for only the Manager machine. Virtualization hosts require read and write access to the ISO domain in order to attach the domain to a data center. If network or host details are not available at the time of setup, or you need to update the ACL at any time, see [Section 12.5, "Changing the Permissions for the Local ISO Domain"](#).

- » The display name, which will be used to label the domain in the Red Hat Enterprise Virtualization Manager. The default name is **ISO_DOMAIN**.

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 ==--

Application mode                : both
Firewall manager                : iptables
Update Firewall                 : True
Host FQDN                       : Your Manager's FQDN
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
Engine installation              : True
NFS setup                       : True
PKI organization                : Your Org
NFS mount point                 : /var/lib/exports/iso
NFS export ACL                  : localhost(rw)
Configure local Engine database : True
Set application as default page  : True
Configure Apache SSL            : True
Configure WebSocket Proxy       : True
Engine Host FQDN                : Your Manager's FQDN

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>.

2.3. Subscribing to the Required Entitlements

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

You can register your system either from the command line or the Subscription Manager GUI.

To install the Red Hat Enterprise Virtualization Manager on a system that does not have access to the Content Delivery Network, see [Section A.1, “Configuring a Local Repository for Offline Red Hat Enterprise Virtualization Manager Installation”](#).



Note

For more information on using the **subscription-manager-gui** utility to register your system from the GUI, see https://access.redhat.com/documentation/en-US/Red_Hat_Subscription_Management/1/html-single/RHSM/index.html#registering-ui.

Procedure 2.1. Subscribing to the Red Hat Enterprise Virtualization Manager entitlements

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

```
# subscription-manager register
```

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

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

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

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



Note

To find out what subscriptions are currently attached, run:

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

To list all enabled repositories, run:

```
# yum repolist
```

4. Disable all existing repositories:

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

5. Enable the required repositories:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-rpms
```

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

2.4. Installing the Red Hat Enterprise Virtualization Manager

2.4.1. 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*m package and dependencies.

Procedure 2.2. 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*m package and dependencies.

```
# yum install rhvm
```



Note

The *rhev*m-*doc* package is installed as a dependency of the *rhev*m 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 command to search for translated versions of the documentation:

```
# yum search rhvm-doc
```

Result

You have installed the *rhev*m package and dependencies.

2.4.2. Preparing a Remote PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager

Optionally configure a PostgreSQL database on a remote Red Hat Enterprise Linux 6.6 machine to use as the Manager database. By default, the Red Hat Enterprise Virtualization Manager's configuration script, **engine-setup**, creates and configures the Manager database locally on the Manager machine. For automatic database configuration, see [Section 2.4.4, "Configuring the Red Hat Enterprise Virtualization Manager"](#). To set up the Manager database with custom values on the Manager machine, see [Section 2.4.3, "Preparing a Local Manually-Configured PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager"](#).

Use this procedure to configure the database on a machine that is separate from the machine where the Manager is installed. Set up this database before you configure the Manager; you must supply the database credentials during **engine-setup**.



Important

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

Procedure 2.3. Preparing a Remote PostgreSQL Database for use with the Red Hat Enterprise Virtualization Manager

1. Install the PostgreSQL server package:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
listen_addresses='*'
```

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

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

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

- Restart the **postgresql** service:

```
# service postgresql restart
```

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

2.4.3. Preparing a Local Manually-Configured PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager

Optionally configure a local PostgreSQL database on the Manager machine to use as the Manager database. By default, the Red Hat Enterprise Virtualization Manager's configuration script, **engine-setup**, creates and configures the Manager database locally on the Manager machine. For automatic database configuration, see [Section 2.4.4, “Configuring the Red Hat Enterprise Virtualization Manager”](#). To configure the Manager database on a machine that is separate from the machine where the Manager is installed, see [Section 2.4.2, “Preparing a Remote PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager”](#).

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



Important

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

Procedure 2.4. Preparing a Local Manually-Configured PostgreSQL Database for use with the Red Hat Enterprise Virtualization Manager

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

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

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

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

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

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

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

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

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

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

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

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

7. Restart the **postgresql** service:

```
# service postgresql restart
```

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

2.4.4. 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.

By default, **engine-setup** creates and configures the Manager database locally on the Manager machine. Alternatively, you can configure the Manager to use a remote database or a manually-configured local database; however, you must set up that database before running **engine-setup**. To set up a remote database see [Section 2.4.2, “Preparing a Remote PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager”](#). To set up a manually-configured local database, see [Section 2.4.3, “Preparing a Local Manually-Configured PostgreSQL Database for Use with the Red Hat Enterprise Virtualization Manager”](#).



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 2.5. Configuring the Red Hat Enterprise Virtualization Manager

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

```
# engine-setup
```

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

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

3. Optionally allow **engine-setup** to configure 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]:
```

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

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

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

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

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

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

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

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

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

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

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

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

```
Database host [localhost]:
Database port [5432]:
Database secured connection (Yes, No) [No]:
Database name [engine]:
Database user [engine]:
Database password:
```

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

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

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

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

Both offers the greatest flexibility.

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

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

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

```
Setup can configure the default page of the web server to present
the application home page. This may conflict with existing
applications.
Do you wish to set the application as the default web page of the
server? (Yes, No) [Yes]:
```

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

```
Setup can configure apache to use SSL using a certificate issued
from the internal CA.
Do you wish Setup to configure that, or prefer to perform that
manually? (Automatic, Manual) [Automatic]:
```

12. Optionally create an NFS share on the Manager to use as an ISO storage domain. The local ISO domain provides a selection of images that can be used in the initial setup of virtual machines:

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

b. Specify the path for the ISO domain:

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

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

`Local ISO domain ACL - note that the default will restrict access to localhost only, for security reasons
[localhost(rw)]: 10.1.2.0/255.255.255.0(rw)
host01.example.com(rw) host02.example.com(rw)`

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

d. Specify a display name for the ISO domain:

`Local ISO domain name [ISO_DOMAIN]:`

13. Optionally, 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]:`

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

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

When your environment has been configured, the **engine-setup** command displays details about how to access your environment. If you chose to manually configure the firewall, **engine-setup** provides a custom list of ports that need to be opened, based on the options selected during setup. The **engine-setup** command also saves your answers to a file that can be used to reconfigure the Manager using the same values, and outputs the location of the log file for the Red Hat Enterprise Virtualization Manager configuration process.

Log in to the Administration Portal as the **admin@internal** user to continue configuring the Manager.

2.4.5. Connecting to the Administration Portal

Access the Administration Portal using a web browser.

Procedure 2.6. Connecting to the Administration Portal

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



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. You must accept this certificate.

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

2.4.6. Removing the Red Hat Enterprise Virtualization Manager

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



Note

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

Procedure 2.7. Removing the 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 whether to remove all Red Hat Enterprise Virtualization Manager components:

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

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

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

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

- 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
```

- Remove the Red Hat Enterprise Virtualization packages:

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

2.4.7. Deploying RHEV-M Virtual Appliance with Self-Hosted Engine

With the RHEV-M Virtual Appliance, you can now use this pre-installed and partially pre-configured image of Red Hat Enterprise Virtualization Manager with your self-hosted engine deployment. The image is available for download as an OVA file from the Customer Portal.



Important

While the packages required to initially set up and run the Manager are pre-loaded into the RHEV-M Virtual Appliance, you must subscribe the operating system in the appliance to the Content Delivery Network to receive support and product updates. For information on subscribing to the required channels, see [Section 3.2, “Subscribing to the Required Entitlements”](#).

Table 2.1. Hardware Requirements

Resource	Minimum	Recommended
Memory	4 GB RAM	16 GB RAM
Disk Space	55 GB writable disk space	100 GB writable disk space

Procedure 2.8. Deploying RHEV-M Virtual Appliance with Self-Hosted Engine

- Log in to the [Customer Portal](#). Open the RHEV-M Virtual Appliance download page at https://access.redhat.com/downloads/content/150/ver=3.5/rhel--6/3.5/x86_64/product-software, and click **Download Now** next to **RHEV-M Appliance (for RHEV-M 3.5)**.
- See [Section 3.3, “Installing the Self-Hosted Engine”](#) for the installation entitlements for a self-hosted engine.
- See [Section 3.4, “Configuring the Self-Hosted Engine”](#) on how to configure a self-hosted engine. In step 4, configuring a virtual machine to be the Red Hat Enterprise Virtualization Manager, specify the option **disk** and the path to the RHEV-M Virtual Appliance file.

```
Please specify the device to boot the VM from (cdrom, disk, pxe)
[cdrom]: disk
Please specify path to OVF archive you would like to use
[None]:/path/to/rhevms.ova
[ INFO ] Checking OVF archive content (could take a few minutes
depending on archive size)
...
```

- After the Red Hat Enterprise Virtualization Manager virtual machine has booted, you will see the RHEV-M Virtual Appliance setup utility. Set the root password and change the default authentication and keyboard configuration as necessary. You will not be able to complete RHN registration at this stage as there will be no network connection.



Important

The root password must be set using the RHEV-M Virtual Appliance setup utility.



- Run the following command to complete your Red Hat Enterprise Virtualization Manager setup:

```
# engine-setup --offline --config-append=rhevms-setup-answers
```

You now have a Red Hat Enterprise Virtualization Manager. To finalize the **hosted-engine** deployment, return to Step 12 "Synchronizing the Host and the Manager" of the procedure in [Section 3.4, "Configuring the Self-Hosted Engine"](#).



Important

SSH is not enabled by default on the RHEV-M Virtual Appliance. You can enable SSH by accessing the Red Hat Enterprise Virtualization Manager virtual machine through the SPICE or VNC console. Edit `/etc/ssh/sshd_config` and change the following two options to **yes**:

- ✧ **PasswordAuthentication**
- ✧ **PermitRootLogin**

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 2.2.

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	No	Yes	No	No

Client Operating System	Wan Optimizations	Dynamic Console Resizing	SPICE Proxy Support	Full High Definition Display	Multiple Monitor Support
Windows Server 2008	Yes	Yes	Yes	Yes	Yes
Windows Server 2012	Yes	No	Yes	No	No

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 can run on Red Hat Enterprise Linux hosts and Red Hat Enterprise Virtualization Hypervisors. Older versions of Red Hat Enterprise Linux are not recommended for use with the self-hosted engine.

This chapter contains information about running a self-hosted engine on Red Hat Enterprise Linux hosts. For information on running a self-hosted engine on Red Hat Enterprise Virtualization Hypervisors, see [Section 7.6.15, “The Hosted Engine Screen”](#).

3.2. Subscribing to the Required Entitlements

Install a Red Hat Enterprise Linux 6.5, 6.6, 6.7, or 7 system. Register the system and subscribe to the required entitlements.

Procedure 3.1. Subscribing to Required Entitlements Using Subscription Manager

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

```
# subscription-manager register
```

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

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

3. 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:

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

4. Disable all existing repositories:

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

5. Enable the required repositories:

A. Red Hat Enterprise Linux 6:

**Important**

For Red Hat Enterprise Linux 6, the *ovirt-hosted-engine-setup* package is provided by the *rhel-6-server-rhev-m-3.5-rpms* repository. If you only have one Red Hat Enterprise Virtualization entitlement, you will need to attach the entitlement to the host first. After you have downloaded the *ovirt-hosted-engine-setup* package, remove the subscription so you can reattach it to the virtual machine to be used as the Manager. See https://access.redhat.com/documentation/en-US/Red_Hat_Subscription_Management/1/html/RHSM/sub-cli.html for more information on how to remove a single product subscription.

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-optional-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-mgmt-agent-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-m-3.5-rpms
```

B. Red Hat Enterprise Linux 7:

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-supplementary-rpms
# subscription-manager repos --enable=rhel-7-server-rhev-mgmt-agent-rpms
```

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

```
# yum update
```

3.3. Installing the Self-Hosted Engine

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.

Prerequisites:

Ensure that you have completed the following prerequisites:

- ✦ You must have a freshly installed Red Hat Enterprise Linux 6.5, 6.6, 6.7, or 7 system and have subscribed it to the required entitlements.
- ✦ You must have prepared either CD-ROM, disk, or PXE installation media for the Manager operating system installation. The physical CD-ROM drive is not supported. To use the CD-ROM option, you must have an ISO file available. For the disk option, you can download the RHEV-M Virtual Appliance for the Manager installation. The RHEV-M Virtual Appliance can be downloaded at <https://rhn.redhat.com/rhn/software/channel/downloads/Download.do?cid=24821>. Note down the full path to the installation media.

- You must have prepared either NFS or iSCSI storage for your self-hosted engine environment. For more information on preparing storage for your deployment, see the [Storage chapter](#) of the *Red Hat Enterprise Virtualization Administration Guide*.
- You must have a fully qualified domain name prepared for your Manager. Forward and reverse lookup records must both be set in the DNS.
- If you are using the RHEV-M Virtual Appliance, the `/tmp` directory must be at least 50 GB.

Procedure 3.2. 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 update
```

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

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

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

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, 6.6, 6.7, or 7 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.

As of Red Hat Enterprise Virtualization 3.5, you have the option to import the RHEV-M Virtual Appliance as the Red Hat Enterprise Virtualization Manager for your self-hosted engine environment. This will need to be available for use before deployment of the **hosted-engine** script. See [Section 2.4.7, “Deploying RHEV-M Virtual Appliance with Self-Hosted Engine”](#) for more information.

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



Note

If deploying the **hosted-engine** customization script over a network, it is recommended to use the **screen** window manager. The package is available in the standard Red Hat Enterprise Linux repository. The benefit of using the **screen** command is that it will preserve the session in case of a network or terminal disruption that would otherwise reset the deployment.

```
# yum install screen
```

Procedure 3.3. 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
```

If deploying the **hosted-engine** customization 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.

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

2. 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*]:

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]:
```

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.



Note

Configuring a bonded and vlan-tagged network interface as the management bridge is currently not supported. To work around this issue, see <https://access.redhat.com/solutions/1417783> for more information.

```
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



Note

For more information on importing the RHEV-M Virtual Appliance as the Red Hat Enterprise Virtualization Manager for your self-hosted engine environment, see [Section 2.4.7, “Deploying RHEV-M Virtual Appliance with Self-Hosted Engine”](#).

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, 6.6, or 6.7 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, and subscribe to the appropriate Red Hat Enterprise Virtualization Manager repositories. See [Subscribing to the Required Entitlements](#).

Ensure that the most up-to-date versions of all installed packages are in use, and install the *rhev* packages.

```
# yum update
```

```
# 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 -h 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 Server (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**.

3.5. Installing Additional Hosts to a Self-Hosted Environment

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, 6.6, 6.7, 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 3.4. 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

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 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.

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 *rhevdm-dwh* and *rhevdm-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.

3.7. Upgrading the Self-Hosted Engine

Upgrade your Red Hat Enterprise Virtualization hosted-engine environment from version 3.4 to 3.5.

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 **3.5**. This avoids any version 3.4 hosts from going into a **Non Operational** state.

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

Procedure 3.5. Upgrading the Self-Hosted Engine

1. Access the Administration Portal. Select Host A and put it into maintenance mode by clicking the **Maintenance** button. If Host A is hosting the Manager virtual machine, the virtual machine will be migrated to Host B.
2. Log in to Host A and set the maintenance mode to **local** to disable the high-availability agents and prevent Host A from hosting the virtual machine while it is being upgraded.

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

3. Enable the new repository on Host A.

```
# subscription-manager repos --enable=rhel-6-server-rhevml-3.5-rpms
```

4. Update Host A.

```
# yum update
```

5. Restart VDSM on Host A.

```
# service vdsmd restart
```

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

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

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

7. Disable maintenance mode to reinstate the high-availability agents on Host A.

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

8. Access the Administration Portal. Select Host A and activate it by clicking the **Activate** button.

9. Disable the old repository on Host A.

```
# subscription-manager repos --disable=rhel-6-server-rhevml-3.4-rpms
```

10. When Host A has a status of **Up**, select Host B and put it into maintenance mode by clicking the **Maintenance** button. This will migrate the Manager virtual machine to Host A.

11. Log in to Host B and set the maintenance mode to **local** to disable the high-availability agents and prevent Host B from hosting the virtual machine while it is being upgraded.

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

12. Enable the new repository on Host B.

```
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
```

13. Update Host B.

```
# yum update
```

14. Restart VDSM on Host B.

```
# service vdsmd restart
```

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

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

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

16. Disable maintenance mode to reinstate the high-availability agents on Host B.

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

17. Access the Administration Portal. Select Host B and activate it by clicking the **Activate** button.

18. Disable the old repository on Host B.

```
# subscription-manager repos --disable=rhel-6-server-rhev-3.4-rpms
```

19. Log in to the Manager virtual machine and update the engine as per the instructions in [Section 5.2.4, “Upgrading to Red Hat Enterprise Virtualization Manager 3.5”](#).

20. Access the Administration Portal.

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

- B. Use the **Compatibility Version** drop-down menu to select **3.5**. Click **OK** to save the change and close the window.

3.8. Upgrading Additional Hosts in a Self-Hosted Environment

It is recommended that all hosts in your self-hosted environment are upgraded at the same time. This prevents version 3.4 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 3.6. Upgrading Additional Hosts

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

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

2. Access the Administration Portal. Select the host and put it into maintenance mode by clicking the **Maintenance** button.
3. Enable the new repository on the host.

```
# subscription-manager repos --enable=rhel-6-server-rhevml-3.5-rpms
```

4. Update the host.

```
# yum update
```

5. Restart VDSM on the host.

```
# service vdsmd restart
```

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

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

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

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

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

8. Access the Administration Portal. Select the host and activate it by clicking the **Activate** button.
9. Disable the old repository on the host.

```
# subscription-manager repos --disable=rhel-6-server-rhevml-3.4-rpms
```

3.9. Removing a Host from a Self-Hosted Engine Environment

If you wish to remove a self-hosted engine host from your environment, you will need to place the host into maintenance, disable the HA services, and remove the self-hosted engine configuration file.

Procedure 3.7. Removing a Host from a Self-Hosted Engine Environment

1. In the Administration Portal, click the **Hosts** resource tab. Select the host, and click **Maintenance** to place the host into maintenance mode.
2. Log in to the host and stop the HA services. This action stops the **ovirt-ha-agent** and **ovirt-ha-broker** services.

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

3. Disable the HA services so the services are not started upon a reboot:

A. Red Hat Enterprise Linux 6:

```
# chkconfig ovirt-ha-agent off
# chkconfig ovirt-ha-broker off
```

B. Red Hat Enterprise Linux 7:

```
# systemctl disable ovirt-ha-agent
# systemctl disable ovirt-ha-broker
```

4. Remove the self-hosted engine configuration file:

A. RHEL-based hypervisor:

```
# rm /etc/ovirt-hosted-engine/hosted-engine.conf
```

B. RHEV-H:

```
# rm /etc/ovirt-hosted-engine/hosted-engine.conf
# unpersist /etc/ovirt-hosted-engine/hosted-engine.conf
```

5. In the Administration Portal, select the same host, and click **Remove** to open the **Remove Host(s)** confirmation window. Click **OK**.

3.10. Backing up and Restoring a Self-Hosted Engine Environment

The nature of the self-hosted engine, and the relationship between the hosts and the hosted-engine virtual machine, means that backing up and restoring a self-hosted engine environment requires additional considerations to that of a standard Red Hat Enterprise Virtualization environment. In particular, the hosted-engine hosts remain in the environment at the time of backup, which can result in a failure to synchronize the new host and hosted-engine virtual machine after the environment has been restored.

To address this, it is recommended that one of the hosts be placed into maintenance mode prior to backup, thereby freeing it from a virtual load. This failover host can then be used to deploy the new self-hosted engine.

If a hosted-engine host is carrying a virtual load at the time of backup, then a host with any of the matching identifiers - IP address, FQDN, or name - cannot be used to deploy a restored self-hosted engine. Conflicts in the database will prevent the host from synchronizing with the restored hosted-engine virtual machine. The failover host, however, can be removed from the restored hosted-engine virtual machine prior to synchronization.



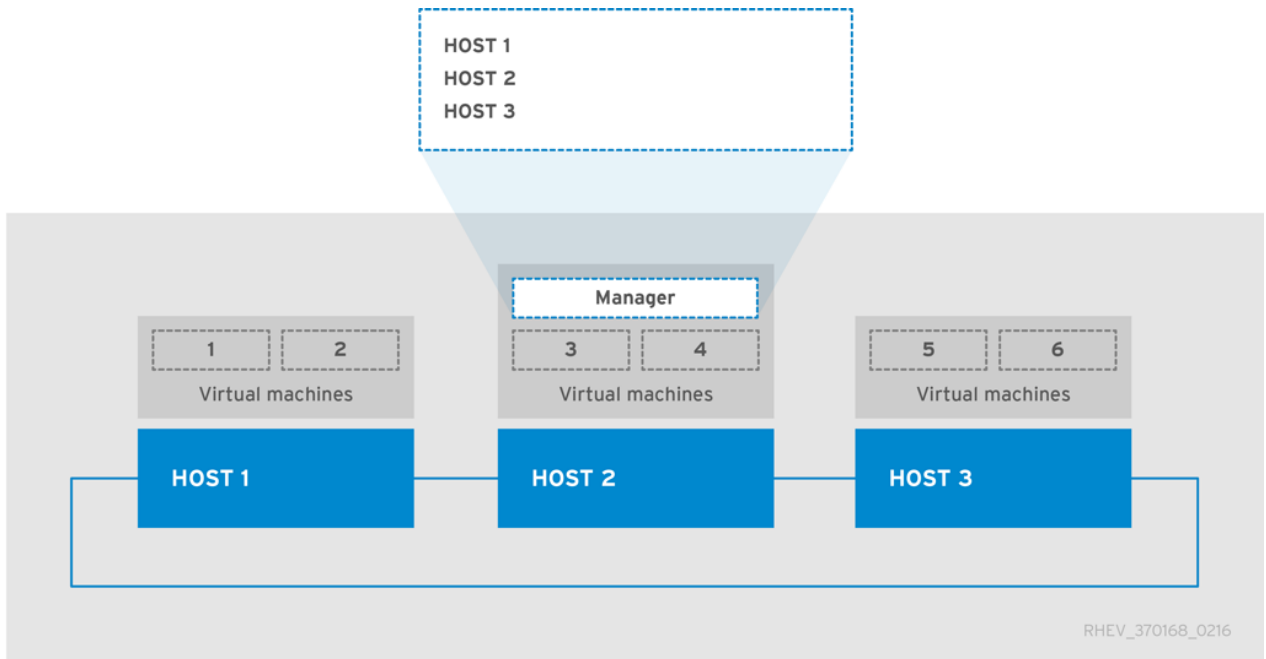
Note

A failover host at the time of backup is not strictly necessary if a new host is used to deploy the self-hosted engine. The new host must have a unique IP address, FQDN, and name so that it does not conflict with any of the hosts present in the database backup.

Procedure 3.8. Workflow for Backing Up the Self-Hosted Engine Environment

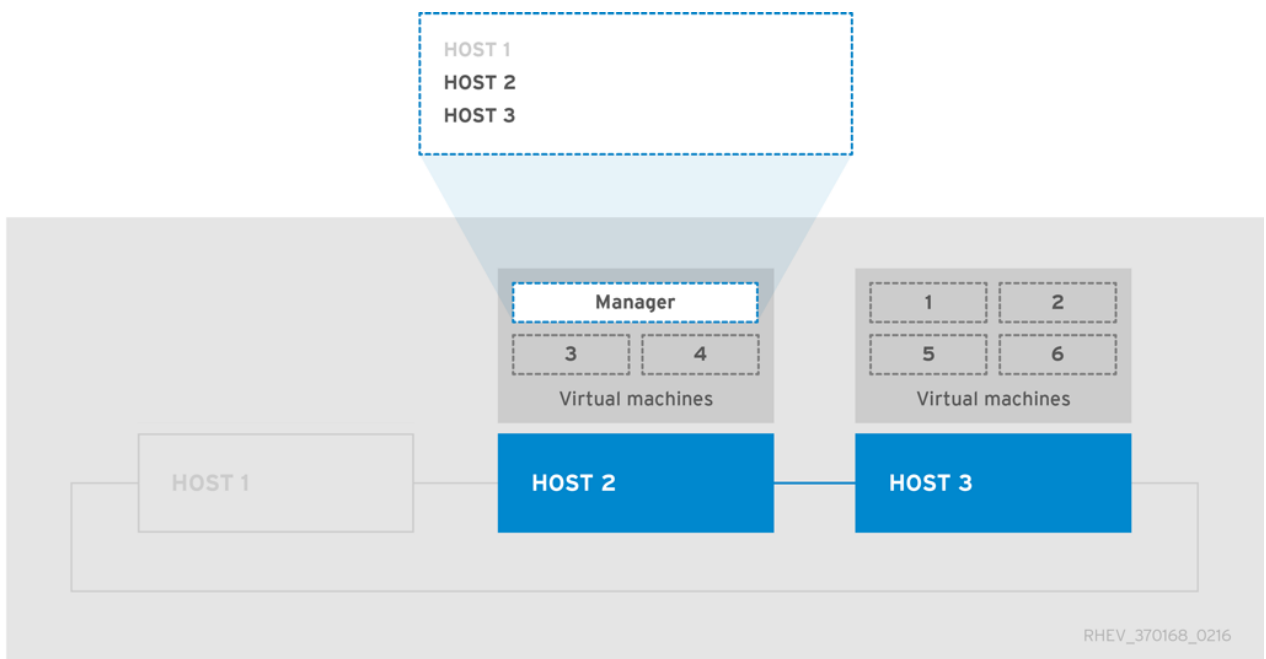
This procedure provides an example of the workflow for backing up a self-hosted engine using a failover host. This host can then be used later to deploy the restored self-hosted engine environment. For more information on backing up the self-hosted engine, see [Section 3.10.1, “Backing up the Self-Hosted Engine Manager Virtual Machine”](#).

1. The engine virtual machine is running on **Host 2** and the six regular virtual machines in the environment are balanced across the three hosts.



Place **Host 1** into maintenance mode. This will migrate the virtual machines on **Host 1** to the other hosts, freeing it of any virtual load and enabling it to be used as a failover host for the backup.

2. **Host 1** is in maintenance mode. The two virtual machines it previously hosted have been migrated to Host 3.

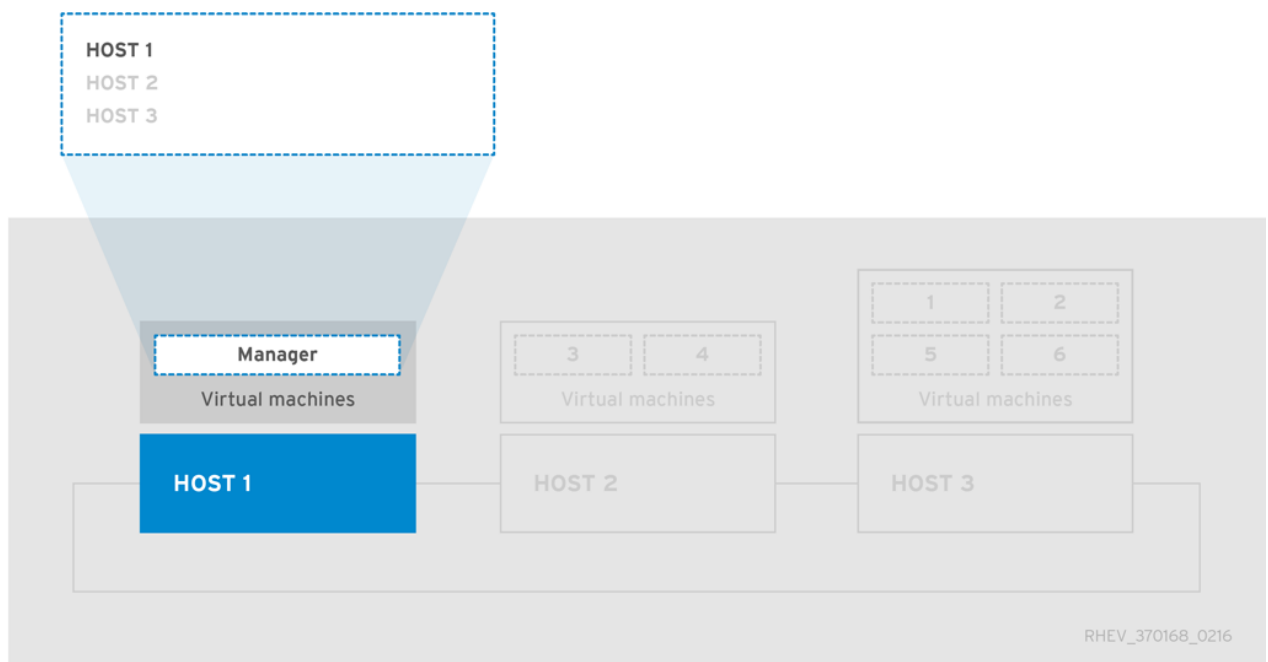


Use **engine-backup** to create backups of the environment. After the backup has been taken, **Host 1** can be activated again to host virtual machines, including the engine virtual machine.

Procedure 3.9. Workflow for Restoring the Self-Hosted Engine Environment

This procedure provides an example of the workflow for restoring the self-hosted engine environment from a backup. The failover host deploys the new engine virtual machine, which then restores the backup. Directly after the backup has been restored, the failover host is still present in the Red Hat Enterprise Virtualization Manager because it was in the environment when the backup was created. Removing the old failover host from the Manager enables the new host to synchronize with the engine virtual machine and finalize deployment. For more information on restoring the self-hosted engine, see [Section 3.10.2, “Restoring the Self-Hosted Engine Environment”](#).

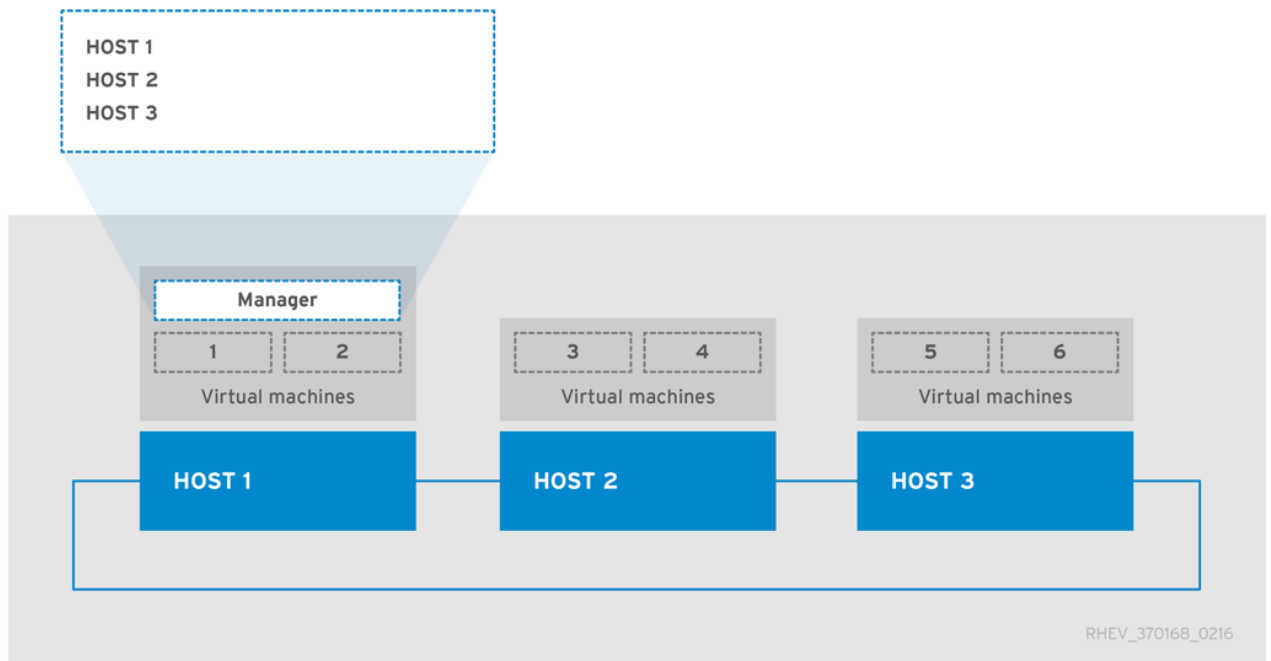
1. **Host 1** has been used to deploy a new self-hosted engine and has restored the backup taken in the previous example procedure. Deploying the restored environment involves additional steps to that of a regular self-hosted engine deployment:
 - ✦ After Red Hat Enterprise Virtualization Manager has been installed on the engine virtual machine, but before **engine-setup** is first run, restore the backup using the **engine-backup** tool.
 - ✦ After **engine-setup** has configured and restored the Manager, log in to the Administration Portal and remove **Host 1**, which will be present from the backup. If old **Host 1** is not removed, and is still present in the Manager when finalizing deployment on new **Host 1**, the engine virtual machine will not be able to synchronize with new **Host 1** and the deployment will fail.



After **Host 1** and the engine virtual machine have synchronized and the deployment has been finalized, the environment can be considered operational on a basic level. With only one hosted-engine host, the engine virtual machine is not highly available. However, if necessary, high-priority virtual machines can be started on **Host 1**.

Any standard RHEL-based hosts - hosts that are present in the environment but are not self-hosted engine hosts - that are operational will become active, and the virtual machines that were active at the time of backup will now be running on these hosts and available in the Manager.

2. **Host 2** and **Host 3** are not recoverable in their current state. These hosts need to be removed from the environment, and then added again to the environment using the hosted-engine deployment script. For more information on these actions, see [Section 3.10.2.3, “Removing Non-Operational Hosts from a Restored Self-Hosted Engine Environment”](#) and [Section 3.10.3, “Installing Additional Hosts to a Restored Self-Hosted Engine Environment”](#).



Host 2 and **Host 3** have been re-deployed into the restored environment. The environment is now as it was in the first image, before the backup was taken, with the exception that the engine virtual machine is hosted on **Host 1**.

3.10.1. 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 3.10. 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

Secure 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.

3.10.2. Restoring the Self-Hosted Engine Environment

This section explains how to restore a self-hosted engine environment from a backup on a newly installed host. The supported restore method uses the **engine-backup** tool.

Restoring a self-hosted engine environment involves the following key actions:

1. Create a freshly installed Red Hat Enterprise Linux host and run the hosted-engine deployment script.
2. Restore the Red Hat Enterprise Virtualization Manager configuration settings and database content in the new Manager virtual machine.
3. 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 newly 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.

3.10.2.1. 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 3.10.1, “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 3.2, “Subscribing to the Required Entitlements”](#) and [Section 3.3, “Installing the Self-Hosted Engine”](#) for more information.

Procedure 3.11. 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. Preparing for Initialization

The script begins by requesting confirmation to use the host as a hypervisor for use in a self-hosted engine environment.

```
Continuing will configure this host for serving as hypervisor and
create a VM where you have to install oVirt Engine afterwards.
Are you sure you want to continue? (Yes, No)[Yes]:
```

4. 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]:
```

5. 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]:
```

6. 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 image alias, 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]:
Please specify an alias for the Hosted Engine image [hosted_engine]:
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]:
```

7. 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]:

8. 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 would like to use.
This needs to match the FQDN that you will use for the engine
installation within the VM.
  Note: This will be the FQDN of the VM you are now going to create,
  it should not point to the base host or to any other existing
  machine.
  Engine FQDN: 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]:
```

9. 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 alias                : hosted_engine
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)[Yes]:

10. 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 ] Stage: Transaction setup
[ INFO ] Stage: Misc configuration
[ INFO ] Stage: Package installation
[ INFO ] Stage: Misc configuration
[ INFO ] Configuring libvirt
[ INFO ] Configuring VDSM
[ INFO ] Starting vdsmd
[ INFO ] Waiting for VDSM hardware info
[ INFO ] Waiting for VDSM hardware info
[ INFO ] Configuring the management bridge
[ INFO ] Creating Storage Domain
[ INFO ] Creating Storage Pool
[ INFO ] Connecting Storage Pool
[ INFO ] Verifying sanlock lockspace initialization
[ 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 "3477XXAM" 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
- (4) Destroy VM and abort setup

(1, 2, 3, 4)[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
```

11. Installing the Virtual Machine Operating System

Connect to Manager virtual machine and install a Red Hat Enterprise Linux 6.5, 6.6, or 6.7 operating system.

12. 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 "3477XXAM" 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 rhevm-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
  (4) Destroy VM and abort setup

(1, 2, 3, 4)[1]:
```

13. Installing the Manager

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

```
# yum upgrade
```

```
# yum install rhelv
```

14. Install Reports and the Data Warehouse

If you are also restoring Reports and the Data Warehouse, install the *rhev*m-*reports-setup* and *rhev*m-*dwh-setup* packages.

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

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

3.10.2.2. 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 3.12. 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.
 - a. 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 *rhev*m package.

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

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

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

- c. Enter the postgresql command line:

```
# su postgres
$ psql
```

- d. 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';
```

e. 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';
```

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

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

g. Edit the **/var/lib/pgsql/data/pg_hba.conf** file as follows:

A. 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
```

B. 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
```

- 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 in [Section 3.10.1, “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 path 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
```

- B. 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 minutes 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.



Note

If the Manager database is restored successfully, but the Manager virtual machine appears to be **Down** and cannot be migrated to another self-hosted engine host, you can enable a new Manager virtual machine and remove the dead Manager virtual machine from the environment by following the steps provided in <https://access.redhat.com/solutions/1517683>.

3.10.2.3. 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 includes the Manager CA certificate file, and uses the **admin@internal** user 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.

3.10.3. 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, 6.7, 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 3.13. 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. Preparing for Initialization

The script begins by requesting confirmation to use the host as a hypervisor for use in a self-hosted engine environment.

```
Continuing will configure this host for serving as hypervisor and  
create a VM where you have to install oVirt Engine afterwards.  
Are you sure you want to continue? (Yes, No)[Yes]:
```

4. 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*]:

5. 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]:

6. 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

7. 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:

8. 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]:

9. 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
(4) Destroy VM and abort setup

(1, 2, 3, 4)[1]

```

10.

```
[ 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
```

11. 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.

3.11. Migrating to a Self-Hosted Environment

To migrate an existing instance of a standard Red Hat Enterprise Virtualization to a self-hosted engine environment, use the **hosted-engine** script to assist with the task. The script asks you a series of questions, and configures your environment based on your answers. The Manager from the standard Red Hat Enterprise Virtualization environment is referred to as the BareMetal-Manager in the following procedure.

The migration involves the following key actions:

- Run the **hosted-engine** script to configure the host to be used as a self-hosted engine host and to create a new Red Hat Enterprise Virtualization virtual machine.
- Back up the the engine database and configuration files using the **engine-backup** tool, copy the backup to the new Manager virtual machine, and restore the backup using the **--mode=restore** parameter of **engine-backup**. Run **engine-setup** to complete the Manager virtual machine configuration.
- Follow the **hosted-engine** script to complete the setup.

Prerequisites

- Prepare a new hypervisor host with the *ovirt-hosted-engine-setup* package installed. See [Section 3.2, “Subscribing to the Required Entitlements”](#) for more information on subscriptions and package installation. The host must be a supported version of the current Red Hat Enterprise Virtualization environment.



Note

If you intend to use an existing host, place the host in maintenance and remove it from the existing environment. See [Administration Guide, Removing a Host](#) for more information.

- Prepare an installation media of the same version of the operating system used for the BareMetal-Manager.
- The fully qualified domain name of the new Manager must be the same fully qualified domain name as that of the BareMetal-Manager. Forward and reverse lookup records must both be set in DNS.

- ✦ You must have access and can make changes to the BareMetal-Manager.

Procedure 3.14. Migrating to a Self-Hosted Environment

1. Initiating Hosted Engine Deployment

Run the **hosted-engine** script. It is recommended to use the **screen** window manager to run the script to avoid losing the session in case of network or terminal disruption. If not already installed, install the **screen** package, which is available in the standard Red Hat Enterprise Linux repository. To escape the script at any time, use the **CTRL+D** keyboard combination to abort deployment.

```
# yum install screen
```

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



Note

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

2. 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. You can only select one iSCSI target during the deployment:

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]:
```

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, 6.6, or 6.7 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 rhevm-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 *rhevm* 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
# 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

See [Restoring the Self-Hosted Engine Manager](#). Skip the **Copying the Backup Files to the New Manager** step in the restore procedure as you have done so in the previous step. After running **engine-setup**, return to this procedure.

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 -h 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

```

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.

3.12. Migrating the Self-Hosted Engine Database to a Remote Server Database

You can migrate the **engine** database of a self-hosted engine to a remote database server after the Red Hat Enterprise Virtualization Manager has been initially configured.

This task is split into two procedures. The first procedure, preparing the remote PostgreSQL database, is a necessary prerequisite for the migration itself and presumes that the server has Red Hat Enterprise Linux installed and has been configured with the appropriate subscriptions.

The second procedure, migrating the database, uses PostgreSQL **pg_dump** and **pg_restore** commands to handle the database backup and restore. As such, it is necessary to edit the **/etc/ovirt-engine/engine.conf.d/10-setup-database.conf** file with the updated information. At a minimum, you must update the location of the new database server. If the database name, role name, or password are modified for the new database server, these values must also be updated in the **10-setup-database.conf** file. This procedure uses the default **engine** database settings to minimize modification of this file.



Note

The Data Warehouse **10-setup-database.conf** file also uses the address of the **engine** database. If Data Warehouse is installed, update the **engine** database values in both the **/etc/ovirt-engine/engine.conf.d/10-setup-database.conf** and the **/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf** file.

Procedure 3.15. Preparing the Remote PostgreSQL Database for use with the Red Hat Enterprise Virtualization Manager

1. Log in to the remote database server and install the PostgreSQL server package:

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

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

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

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

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

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

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

**Note**

The password for the **engine** user is located in plain text in `/etc/ovirt-engine/engine.conf.d/10-setup-database.conf`. Any password can be used when creating the role on the new server, however if a different password is used then this file must be updated with the new password.

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

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

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

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

7. 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.

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

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

9. Restart the **postgresql** service:

```
# service postgresql restart
```

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

Procedure 3.16. Migrating the Database

1. Log in to one of the hosted-engine hosts and place the environment into **global** maintenance mode so that the High Availability agents do not interfere with the Manager virtual machine during the database migration:

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

2. Log in to the Manager virtual machine and stop the **ovirt-engine** service so that it does not interfere with the engine backup:

```
# service ovirt-engine stop
```

3. Create the **engine** database backup using the PostgreSQL **pg_dump** command:

```
# su - postgres -c 'pg_dump -F c engine -f /tmp/engine.dump'
```

4. Copy the backup file to the new database server. The target directory must allow write access for the **postgres** user:

```
# scp /tmp/engine.dump  
root@new.database.server.com:/tmp/engine.dump
```

5. Log in to the new database server and restore the database using the PostgreSQL **pg_restore** command:

```
# su - postgres -c 'pg_restore -d engine /tmp/engine.dump'
```

6. Log in to the Manager virtual machine and update the **/etc/ovirt-engine/engine.conf.d/10-setup-database.conf** and replace the **localhost** value of **ENGINE_DB_HOST** with the IP address of the new database server. If the engine name, role name, or password differ on the new database server, update those values in this file.

If Data Warehouse is installed, these values also need to be updated in the **/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf** file.

7. Now that the database has been migrated, start the **ovirt-engine** service:

```
# service ovirt-engine start
```

8. Log in to a hosted-engine host and turn off maintenance mode, enabling the High Availability agents:

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

3.13. Migrating the Data Warehouse Database to a Remote Server Database

You can migrate the **ovirt_engine_history** database to a remote database server after the Red Hat Enterprise Virtualization Manager has been initially configured.

This task is split into two procedures. The first procedure, preparing the remote PostgreSQL database, is a necessary prerequisite for the migration itself and presumes that the server has Red Hat Enterprise Linux installed and has been configured with the appropriate subscriptions.

The second procedure, migrating the database, uses PostgreSQL **pg_dump** and **pg_restore** commands to handle the database backup and restore. As such, it is necessary to edit the **/etc/ovirt-engine-reports/ovirt-engine-reports.conf.d/10-setup-database.conf** and **/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf** files with the updated information. At a minimum, you must update the location of

the new database server. If the database name, role name, or password are modified for the new database server, these values must also be updated in both **10-setup-database.conf** files. This procedure uses the default **ovirt_engine_history** database settings to minimize modification of this file.

Procedure 3.17. Preparing the Remote PostgreSQL Database for use with the Red Hat Enterprise Virtualization Manager

1. Log in to the remote database server and install the PostgreSQL server package:

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

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

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

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

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

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

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



Note

The password for the **ovirt_engine_history** user is located in plain text in **/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf**. Any password can be used when creating the role on the new server, however if a different password is used then this file, and the **/etc/ovirt-engine-reports/ovirt-engine-reports.conf.d/10-setup-database.conf** file, must be updated with the new password.

5. Create a database in which to store the history of the Red Hat Enterprise Virtualization environment. The default database name is **ovirt_engine_history**, and the default user name is **ovirt_engine_history**:

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

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

```
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 --dport 5432 -j ACCEPT
# service iptables save
```

- Restart the **postgresql** service:

```
# service postgresql restart
```

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

Procedure 3.18. Migrating the Database

- Log in to one of the hosted-engine hosts and place the environment into **global** maintenance mode so that the High Availability agents do not interfere with the Manager virtual machine during the database migration:

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

- Log in to the Red Hat Enterprise Virtualization Manager machine and stop the **ovirt-engine-dwhd** service so that it does not interfere with the engine backup:

```
# service ovirt-engine-dwhd stop
```

- Create the **ovirt_engine_history** database backup using the PostgreSQL **pg_dump** command:

```
# su - postgres -c 'pg_dump -F c ovirt_engine_history -f /tmp/ovirt_engine_history.dump'
```

- Copy the backup file to the new database server. The target directory must allow write access for the **postgres** user:

```
# scp /tmp/ovirt_engine_history.dump
root@new.database.server.com:/tmp/ovirt_engine_history.dump
```

- Log in to the new database server and restore the database using the PostgreSQL **pg_restore** command:

```
# su - postgres -c 'pg_restore -d ovirt_engine_history /tmp/ovirt_engine_history.dump'
```


- Log in to the Manager server and update the `/etc/ovirt-engine-reports/ovirt-engine-reports.conf.d/10-setup-database.conf` and `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` files, replacing the `localhost` value of `DWH_DB_HOST` with the IP address of the new database server. If the `DWH_DB_DATABASE`, `DWH_DB_USER`, or `DWH_DB_PASSWORD` differ on the new database server, update those values in these files.

If the Manager database has also been migrated, these values must also be updated in the `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file.

- Use a web browser to log in to the Reports portal at

`https://hostname.example.com/ovirt-engine-reports`

using the **superuser** user name. Click **View** → **Repository** to open the **Folders** side pane.

- In the **Folders** side pane, select **RHEVM Reports** → **Resources** → **JDBC** → **Data Sources**.

- Select **oVirt History** and click **Edit**.

- Update the **Host (required)** field with the IP address of the new database server and click **Save**.

- Now that the database has been migrated and the Reports portal connects to it, start the **ovirt-engine-dwhd** service:

```
# service ovirt-engine-dwhd start
```

- Log in to a hosted-engine host and turn off maintenance mode, enabling the High Availability agents:

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

3.14. Migrating the Reports Database to a Remote Server Database

You can migrate the **ovirt_engine_reports** database to a remote database server after the Red Hat Enterprise Virtualization Manager has been initially configured.

This task is split into two procedures. The first procedure, preparing the remote PostgreSQL database, is a necessary prerequisite for the migration itself and presumes that the server has Red Hat Enterprise Linux installed and has been configured with the appropriate subscriptions.

The second procedure, migrating the database, uses PostgreSQL **pg_dump** and **pg_restore** commands to handle the database backup and restore. As such, it is necessary to edit the `/var/lib/ovirt-engine-reports/build-conf/master.properties` file with the updated information. At a minimum, you must update the location of the new database server. If the database name, role name, or password are modified for the new database server, these values must also be updated in both **master.properties** files. This procedure uses the default **ovirt_engine_reports** database settings to minimize modification of this file.

Procedure 3.19. Preparing the Remote PostgreSQL Database for use with the Red Hat Enterprise Virtualization Manager

1. Log in to the remote database server and install the PostgreSQL server package:

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

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

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

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

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

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

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



Note

The password for the **ovirt_engine_reports** user is located in plain text in **/var/lib/ovirt-engine-reports/build-conf/master.properties**. Any password can be used when creating the role on the new server, however if a different password is used then this file must be updated with the new password.

5. Create a database in which to store the history of the Red Hat Enterprise Virtualization environment. The default database name is **ovirt_engine_reports**, and the default user name is **ovirt_engine_reports**:

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

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

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

7. 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 --dport 5432 -j ACCEPT
# service iptables save
```

- Restart the **postgresql** service:

```
# service postgresql restart
```

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

Procedure 3.20. Migrating the Database

- Log in to one of the hosted-engine hosts and place the environment into **global** maintenance mode so that the High Availability agents do not interfere with the Manager virtual machine during the database migration:

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

- Log in to the Red Hat Enterprise Virtualization Manager machine and stop the **ovirt-engine-reportsd** service so that it does not interfere with the engine backup:

```
# service ovirt-engine-reportsd stop
```

- Create the **ovirt_engine_reports** database backup using the PostgreSQL **pg_dump** command:

```
# su - postgres -c 'pg_dump -F c ovirt_engine_reports -f /tmp/ovirt_engine_reports.dump'
```

- Copy the backup file to the new database server. The target directory must allow write access for the **postgres** user:

```
# scp /tmp/ovirt_engine_reports.dump
root@new.database.server.com:/tmp/ovirt_engine_reports.dump
```

- Log in to the new database server and restore the database using the PostgreSQL **pg_restore** command:

```
# su - postgres -c 'pg_restore -d ovirt_engine_reports /tmp/ovirt_engine_reports.dump'
```

- Log in to the Manager server and update **/var/lib/ovirt-engine-reports/build-conf/master.properties**, replacing the **localhost** value of **dbHost** with the IP address of the new database server. If the **ovirt_engine_reports.js.dbName**, **dbUsername**, or **dbPassword** differ on the new database server, update those values in this file.
- Now that the database has been migrated, you must run **engine-setup** to rebuild reports with the new credentials:

```
# engine-setup
```

8. Log in to a hosted-engine host and turn off maintenance mode, enabling the High Availability agents:

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

Chapter 4. Data Warehouse and Reports

4.1. Workflow Progress - Data Collection Setup and Reports Installation



4.2. Overview of Configuring Data Warehouse and Reports

The Red Hat Enterprise Virtualization Manager 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. Installing Data Warehouse creates the **ovirt_engine_history** database, to which the Manager is configured to log information for reporting purposes. Red Hat Enterprise Virtualization Manager Reports functionality is also available as an optional component. Reports provides a customized implementation of JasperServer and JasperReports, an open source reporting tool capable of being embedded in Java-based applications. It produces reports that can be built and accessed via a web user interface, and then rendered to screen, printed, or exported to a variety of formats including PDF, Excel, CSV, Word, RTF, Flash, ODT and ODS. The Data Warehouse and Reports components are optional, and must be installed and configured in addition to the Manager setup.

Before proceeding with Data Warehouse and Reports installation you must first have installed and configured the Red Hat Enterprise Virtualization Manager. The Reports functionality depends on the presence of the Data Warehouse; Data Warehouse must be installed and configured before Reports.

It is recommended that you set the system time zone for all machines in your Data Warehouse/Reports deployment to UTC. This ensures that data collection is not interrupted by variations in your local time zone: for example, a change from summer time to winter time.

To calculate an estimate of the space and resources the **ovirt_engine_history** database will use, use the [RHEV Manager History Database Size Calculator](#) tool. The estimate is based on the number of entities and the length of time you have chosen to retain the history records.

4.3. Data Warehouse and Reports Configuration Notes

Behavior

The following behavior is expected in **engine-setup**:

Install the Data Warehouse package and the Reports package, run **engine-setup**, and answer **No** to configuring Data Warehouse and Reports:

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

Run **engine-setup** again; setup no longer presents the option to configure those services.

Workaround

To force **engine-setup** to present both options again, run **engine-setup** with the following options appended:

```
# engine-setup --otopi-
environment='OVESETUP_REPORTS_CORE/enable=none:None
OVESETUP_DWH_CORE/enable=none:None'
```

To present only the Data Warehouse option, run:

```
# engine-setup --otopi-environment='OVESETUP_DWH_CORE/enable=none:None'
```

To present only the Reports option, run:

```
# engine-setup --otopi-
environment='OVESETUP_REPORTS_CORE/enable=none:None'
```



Note

To configure only the currently installed Data Warehouse and Reports packages, and prevent setup from applying package updates found in enabled repositories, add the **--offline** option .

4.4. Data Warehouse and Reports Installation Options

Data Warehouse and Reports installation requires between one and three machines, and can be configured in one of the following ways:

1. **Install and configure both Data Warehouse and Reports on the machine on which the Manager is installed.**

This configuration hosts the Data Warehouse and Reports services on your Manager machine. This requires only a single registered machine, and is the simplest to configure; however, it also requires that the services share CPU and memory, and increases the demand on the host machine. Users who require access to the Data Warehouse service or the Reports service will require access to the Manager machine itself.

2. **Install and configure both Data Warehouse and Reports on one separate machine.**

This configuration hosts Data Warehouse and Reports on a single, separate machine. This requires two registered machines; however, it reduces the load on the Manager machine, and avoids potential CPU and memory-sharing conflicts on that machine. Administrators can also allow user access to the Data Warehouse-Reports machine, without the need to grant access to the Manager machine. Note that the Data Warehouse and Reports services will still compete for resources on their single host.

3. **Install and configure Data Warehouse on a separate machine, then install and configure Reports on a separate machine.**

This configuration separates each service onto its own dedicated host. This requires three registered machines; however, it reduces the load on each individual machine, and allows each service to avoid potential conflicts caused by sharing CPU and memory with other processes. Administrators can also allow user access to one particular machine, without the need to grant access to either of the two other machines.

4. **Install and configure Data Warehouse on the Manager machine, then install and configure Reports on a separate machine.**

This configuration hosts Data Warehouse on the Manager machine, and Reports on a separate host. This requires two registered machines; however, it reduces the load on the Manager machine, and avoids some memory-sharing conflicts. Administrators can allow user access to the Reports machine, without the need to grant access to the Manager machine.

5. **Install and configure Data Warehouse on a separate machine, then install and configure Reports on the Manager machine.**

This configuration hosts Data Warehouse on a separate machine, and Reports on the Manager machine. This requires two registered machines; however, it reduces the load on the Manager machine, and avoids some memory-sharing conflicts. Administrators can allow user access to the Data Warehouse machine, without the need to grant access to the Manager machine.

If you choose to host the Data Warehouse database on a machine that is separate from the machine on which the Data Warehouse service is installed, you will require an additional machine for that purpose. The same is true if you choose to host the Reports database remotely.

**Note**

Detailed user, administration, and installation guides for JasperReports are available in `/usr/share/jasperreports-server-pro/docs/`

4.4.1. Installing and Configuring Data Warehouse and Reports on the Red Hat Enterprise Virtualization Manager

Overview

Install and configure Data Warehouse and Red Hat Enterprise Virtualization Manager Reports on the same machine as the Red Hat Enterprise Virtualization Manager.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager on this machine.
2. If you choose to use a remote Data Warehouse database or Reports database, you must set up each database before installing the Data Warehouse and Reports services. You must have the following information about each database host:
 - ✳ The fully qualified domain name of the host
 - ✳ The port through which the database can be reached (5432 by default)
 - ✳ The database name
 - ✳ The database user
 - ✳ The database password
3. 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 Data Warehouse and Reports on the Red Hat Enterprise Virtualization Manager

1. Install the *rhev-dwh* package and the *rhev-reports* package on the system where the Red Hat Enterprise Virtualization Manager is installed:

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

2. Run the **engine-setup** command to begin configuration of Data Warehouse and Reports on the machine:

```
# engine-setup
```

3. Follow the prompts to configure Data Warehouse and Reports:

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


- Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

- Answer the following questions about the Data Warehouse database and the Reports database:

```
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 and then press **Enter**. If you select **Remote**, you are prompted to provide details about each remote database host.

- Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

```
Reports power users password:
```

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

- For the configuration 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]:
```

Press **Enter** to proceed. The **ovirt-engine** service restarts automatically later in the command.

- Confirm your installation settings:

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

Access the Reports Portal at <http://demo.redhat.com/ovirt-engine-reports>, replacing *demo.redhat.com* with the fully qualified domain name of the Manager. If during the Manager installation you selected a non-default HTTP port then append *:port* to the URL, replacing *:port* with the port that you chose.

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

4.4.2. Installing and Configuring Data Warehouse and Reports on the Same Separate Machine

Overview

Install and configure Data Warehouse and Red Hat Enterprise Virtualization Manager Reports together on a separate host from that on which the Red Hat Enterprise Virtualization Manager is installed. Hosting the Data Warehouse service and the Reports service on a separate machine helps to reduce the load on the Manager machine. Note that hosting Data Warehouse and Reports on the same machine means that these processes will share CPU and memory.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager on a separate machine.
2. To set up the Data Warehouse and Reports machine, you must have the following:
 - ✳ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✳ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** subscription pools.
 - ✳ The password from the Manager's `/etc/ovirt-engine/engine.conf.d/10-setup-database.conf` file.
 - ✳ Allowed access from the Data Warehouse-Reports machine to the Manager database machine's TCP port 5432.
3. If you choose to use a remote Data Warehouse database or Reports database, you must set up each database before installing the Data Warehouse and Reports services. You must have the following information about each database host:
 - ✳ The fully qualified domain name of the host
 - ✳ The port through which the database can be reached (5432 by default)
 - ✳ The database name
 - ✳ The database user
 - ✳ The database password

Procedure 4.2. Installing and Configuring Data Warehouse and Reports on the Same Separate Machine

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

```
# subscription-manager register
```

- Find subscription pools containing the repositories required to install Data Warehouse and Reports:

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

- 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:

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

- Disable all existing repositories:

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

- Enable the required repositories:

```
# subscription-manager repos --enable=rhel-6-server-rpms  
# subscription-manager repos --enable=rhel-6-server-supplementary-  
rpms  
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms  
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-  
rpms
```

- Ensure that all packages currently installed are up to date:

```
# yum update
```

- Install the *rhev-dwh-setup* and *rhev-reports-setup* packages:

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

- Run the **engine-setup** command to begin configuration of Data Warehouse and Reports on the machine:

```
# engine-setup
```

- Follow the prompts to configure Data Warehouse and Reports:

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

- Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

Host fully qualified DNS name of this server [*autodetected hostname*]:

12. Enter the fully qualified domain name of the Manager machine, and then press **Enter**:

Host fully qualified DNS name of the engine server []:

13. Answer the following questions about the Data Warehouse database and the Reports database:

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 and then press **Enter**. If you select **Remote**, you are prompted to provide details about each remote database host.

14. Enter the fully qualified domain name and password for the Manager database machine. Press **Enter** to accept the default values in each other field:

Engine database host []: *engine-db-fqdn*
 Engine database port [5432]:
 Engine database secured connection (Yes, No) [No]:
 Engine database name [engine]:
 Engine database user [engine]:
 Engine database password: *password*

15. Press **Enter** to allow setup to sign the Reports certificate and Apache certificate on the Manager via SSH:

Setup will need to do some actions on the remote engine server.

```

Either automatically, using ssh as root to access it, or you will
be prompted to manually perform each such action.
Please choose one of the following:
1 - Access remote engine server using ssh as root
2 - Perform each action manually, use files to copy content around
(1, 2) [1]:

```

16. Press **Enter** to accept the default SSH port, or enter an alternative port number and then press **Enter**:

```
ssh port on remote engine server [22]:
```

17. Enter the root password for the Manager machine:

```
root password on remote engine server manager-fqdn.com:
```

18. Press **Enter** to allow automatic configuration of SSL on Apache:

```

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]:

```

19. Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

```
Reports power users password:
```

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

20. Confirm your installation settings:

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

Next Steps

Access the Reports Portal at <http://demo.redhat.com/ovirt-engine-reports>, replacing *demo.redhat.com* with the fully qualified domain name of the Manager. If during the Manager installation you selected a non-default HTTP port then append *:port* to the URL, replacing *:port* with the port that you chose.

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

4.4.3. Installing and Configuring Data Warehouse and Reports on Separate Machines

Overview

Install and configure Data Warehouse on a separate host from that on which the Red Hat Enterprise Virtualization Manager is installed, then install and configure Red Hat Enterprise Virtualization Manager Reports on a third machine. Hosting the Data Warehouse and Reports services on separate machines helps to reduce the load on the Manager machine. Separating Data Warehouse and

Reports onto individual machines further reduces the demand each service places on its host machine, and avoids any conflicts caused by sharing CPU and memory with other processes.

Installing this scenario involves two key steps:

1. Install and configure Data Warehouse on a separate machine.
2. Install and configure Reports on a separate machine.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager on a separate machine.
2. To set up the Data Warehouse machine, you must have the following:
 - ✧ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✧ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** subscription pools.
 - ✧ The password from the Manager's `/etc/ovirt-engine/engine.conf.d/10-setup-database.conf` file.
 - ✧ Allowed access from the Data Warehouse machine to the Manager database machine's TCP port 5432.
3. To set up the Reports machine, you must have the following:
 - ✧ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✧ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** subscription pools.
 - ✧ The password from the Data Warehouse machine's `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file.
 - ✧ Allowed access from the Reports machine to the Manager database machine's TCP port 5432.
4. If you choose to use a remote Data Warehouse database or Reports database, you must set up each database before installing the Data Warehouse and Reports services. You must have the following information about each database host:
 - ✧ The fully qualified domain name of the host
 - ✧ The port through which the database can be reached (5432 by default)
 - ✧ The database name
 - ✧ The database user
 - ✧ The database password

Procedure 4.3. Step 1: Installing and Configuring Data Warehouse on a Separate Machine

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

```
# subscription-manager register
```

- Find subscription pools containing the repositories required to install Data Warehouse:

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

- 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:

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

- Disable all existing repositories:

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

- Enable the required repositories:

```
# subscription-manager repos --enable=rhel-6-server-rpms  
# subscription-manager repos --enable=rhel-6-server-supplementary-  
rpms  
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms  
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-  
rpms
```

- Ensure that all packages currently installed are up to date:

```
# yum update
```

- Install the *rhev-dwh-setup* package:

```
# yum install rhel-dwh-setup
```

- Run the **engine-setup** command to begin configuration of Data Warehouse on the machine:

```
# engine-setup
```

- Press **Enter** to configure Data Warehouse:

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

- Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

12. Answer the following questions about the Data Warehouse database:

```
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]:
```

Press **Enter** to choose the highlighted defaults, or type your alternative preference and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

13. Enter the fully qualified domain name and password for the Manager database machine. Press **Enter** to accept the default values in each other field:

```
Engine database host []: engine-db-fqdn
Engine database port [5432]:
Engine database secured connection (Yes, No) [No]:
Engine database name [engine]:
Engine database user [engine]:
Engine database password: password
```

14. Confirm your installation settings:

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

Procedure 4.4. Step 2: Installing and Configuring Reports on a Separate Machine

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

```
# subscription-manager register
```

2. Find subscription pools containing the repositories required to install Reports:

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

3. 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:


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

4. Disable all existing repositories:

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

5. Enable the required repositories:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-rpms
```

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

```
# yum update
```

7. Install the *rhev-reports-setup* package:

```
# yum install rhv-reports-setup
```

8. Run the **engine-setup** command to begin configuration of Reports on the machine:

```
# engine-setup
```

9. Press **Enter** to configure Reports:

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

10. Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

12. Enter the fully qualified domain name of the Manager machine, and then press **Enter**:

Host fully qualified DNS name of the engine server []:

13. Answer the following questions about the Reports database:

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 and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

14. Enter the fully qualified domain name and password for your Data Warehouse database host. Press **Enter** to accept the default values in each other field:

DWH database host []: *dwh-db-fqdn*
 DWH database port [5432]:
 DWH database secured connection (Yes, No) [No]:
 DWH database name [ovirt_engine_history]:
 DWH database user [ovirt_engine_history]:
 DWH database password: *password*

15. Press **Enter** to allow setup to sign the Reports certificate and Apache certificate on the Manager via SSH:

Setup will need to do some actions on the remote engine server. Either automatically, using ssh as root to access it, or you will be prompted to manually perform each such action. Please choose one of the following:
 1 - Access remote engine server using ssh as root
 2 - Perform each action manually, use files to copy content around (1, 2) [1]:

16. Press **Enter** to accept the default SSH port, or enter an alternative port number and then press **Enter**:

ssh port on remote engine server [22]:

17. Enter the root password for the Manager machine:

root password on remote engine server *manager-fqdn.com*:

18. Press **Enter** to allow automatic configuration of SSL on Apache:

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]:

19. Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

Reports power users password:

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

20. Confirm your installation settings:

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

Next Steps

Access the Reports Portal at <http://demo.redhat.com/ovirt-engine-reports>, replacing *demo.redhat.com* with the fully qualified domain name of the Manager. If during the Manager installation you selected a non-default HTTP port then append *:port* to the URL, replacing *:port* with the port that you chose.

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

4.4.4. Installing and Configuring Data Warehouse on the Red Hat Enterprise Virtualization Manager and Reports on a Separate Machine

Overview

Install and configure Data Warehouse on the same system as the Red Hat Enterprise Virtualization Manager, then install and configure Red Hat Enterprise Virtualization Manager Reports on a separate machine. Hosting the Reports service on a separate machine helps to reduce the load on the Manager machine.

Installing this scenario involves two key steps:

1. Install and configure Data Warehouse on the Manager machine.
2. Install and configure Reports on a separate machine.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager on one machine. This is the machine on which you are installing Data Warehouse.
2. To set up the Reports machine, you must have the following:
 - ✦ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✦ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlement pools.
 - ✦ The password from the Data Warehouse machine's `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file.
 - ✦ Allowed access from the Reports machine to the Manager database machine's TCP port 5432.

3. If you choose to use a remote Data Warehouse database or Reports database, you must set up each database before installing the Data Warehouse and Reports services. You must have the following information about each database host:
 - ✳ The fully qualified domain name of the host
 - ✳ The port through which the database can be reached (5432 by default)
 - ✳ The database name
 - ✳ The database user
 - ✳ The database password
4. If you are using the Self-Hosted Engine, you must move it to maintenance mode:

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

Procedure 4.5. Step 1: Installing and Configuring Data Warehouse on the Manager Machine

1. Install the *rhevdm-dwh* package:

```
# yum install rhvdm-dwh
```

2. Run the **engine-setup** command to begin configuration of Data Warehouse on the machine:

```
# engine-setup
```

3. Press **Enter** to configure Data Warehouse:

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

4. Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

5. Answer the following questions about the Data Warehouse database:

```
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]:
```

Press **Enter** to choose the highlighted defaults, or type your alternative preference and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

- For the configuration 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]:
```

Press **Enter** to proceed. The **ovirt-engine** service restarts automatically later in the command.

- Confirm your installation settings:

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

Procedure 4.6. Step 2: Installing and Configuring Reports on a Separate Machine

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

```
# subscription-manager register
```

- Find entitlement pools containing the channels required to install Reports:

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

- 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:

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

- Disable all existing repositories:

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

- Enable the required channels:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-
rpms
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-
rpms
```

- Ensure that all packages currently installed are up to date:

```
# yum update
```

- Install the *rhev-reports-setup* package:

```
# yum install rhelm-reports-setup
```

8. Run the **engine-setup** command to begin configuration of Reports on the machine:

```
# engine-setup
```

9. Press **Enter** to configure Reports:

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

10. Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

12. Enter the fully qualified domain name of the Manager machine, and then press **Enter**:

```
Host fully qualified DNS name of the engine server []:
```

13. Answer the following questions about the Reports database:

```
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 and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

14. Enter the fully qualified domain name and password for your Data Warehouse database host. Press **Enter** to accept the default values in each other field:

```
DWH database host []: dwh-db-fqdn
DWH database port [5432]:
DWH database secured connection (Yes, No) [No]:
DWH database name [ovirt_engine_history]:
```

```
DWH database user [ovirt_engine_history]:
DWH database password: password
```

15. Press **Enter** to allow setup to sign the Reports certificate and Apache certificate on the Manager via SSH:

```
Setup will need to do some actions on the remote engine server.
Either automatically, using ssh as root to access it, or you will
be prompted to manually perform each such action.
Please choose one of the following:
1 - Access remote engine server using ssh as root
2 - Perform each action manually, use files to copy content around
(1, 2) [1]:
```

16. Press **Enter** to accept the default SSH port, or enter an alternative port number and then press **Enter**:

```
ssh port on remote engine server [22]:
```

17. Enter the root password for the Manager machine:

```
root password on remote engine server manager-fqdn.com:
```

18. Press **Enter** to allow automatic configuration of SSL on Apache:

```
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]:
```

19. Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

```
Reports power users password:
```

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

20. Confirm your installation settings:

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

Next Steps

Access the Reports Portal at <http://demo.redhat.com/ovirt-engine-reports>, replacing *demo.redhat.com* with the fully qualified domain name of the Manager. If during the Manager installation you selected a non-default HTTP port then append *:port* to the URL, replacing *:port* with the port that you chose.

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

4.4.5. Installing and Configuring Data Warehouse on a Separate Machine and Reports on the Red Hat Enterprise Virtualization Manager

Reports on the Red Hat Enterprise Virtualization Manager

Overview

Install and configure Data Warehouse on a separate host from that on which the Red Hat Enterprise Virtualization Manager is installed, then install and configure Red Hat Enterprise Virtualization Manager Reports on the Manager machine. Hosting the Data Warehouse service on a separate machine helps to reduce the load on the Manager machine. Note that hosting the Manager and Reports on the same machine means that these processes will share CPU and memory.

Installing this scenario involves two key steps:

1. Install and configure Data Warehouse on a separate machine.
2. Install and configure Reports on the Manager machine.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager on a separate machine.
2. To set up the Data Warehouse machine, you must have the following:
 - ✳ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✳ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlement pools.
 - ✳ The password from the Manager's `/etc/ovirt-engine/engine.conf.d/10-setup-database.conf` file.
 - ✳ Allowed access from the Data Warehouse machine to the Manager database machine's TCP port 5432.
3. To set up the Reports machine, you must have the following:
 - ✳ The password from the Data Warehouse machine's `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file.
4. If you choose to use a remote Data Warehouse database or Reports database, you must set up each database before installing the Data Warehouse and Reports services. You must have the following information about each database host:
 - ✳ The fully qualified domain name of the host
 - ✳ The port through which the database can be reached (5432 by default)
 - ✳ The database name
 - ✳ The database user
 - ✳ The database password
5. If you are using the Self-Hosted Engine, you must move it to maintenance mode:

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

Procedure 4.7. Step 1: Installing and Configuring Data Warehouse on a Separate Machine

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

```
# subscription-manager register
```

2. Find entitlement pools containing the channels required to install Data Warehouse:

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

3. 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:

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

4. Disable all existing repositories:

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

5. Enable the required channels:

```
# subscription-manager repos --enable=rhel-6-server-rpms  
# subscription-manager repos --enable=rhel-6-server-supplementary-  
rpms  
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms  
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-  
rpms
```

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

```
# yum update
```

7. Install the *rhev-dwh-setup* package:

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

8. Run the **engine-setup** command to begin configuration of Data Warehouse on the machine:

```
# engine-setup
```

9. Press **Enter** to configure Data Warehouse:

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

10. Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

12. Answer the following questions about the Data Warehouse database:

```
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]:
```

Press **Enter** to choose the highlighted defaults, or type your alternative preference and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

13. Enter the fully qualified domain name and password for the Manager database machine. Press **Enter** to accept the default values in each other field:

```
Engine database host []: engine-db-fqdn
Engine database port [5432]:
Engine database secured connection (Yes, No) [No]:
Engine database name [engine]:
Engine database user [engine]:
Engine database password: password
```

14. Confirm your installation settings:

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

Procedure 4.8. Step 2: Installing and Configuring Reports on the Manager Machine

1. Install the *rhevdm-reports* package:

```
# yum install rhvdm-reports
```

2. Run the **engine-setup** command to begin configuration of Reports on the machine:

```
# engine-setup
```

- Press **Enter** to configure Reports:

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

- Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

- Answer the following questions about the Reports database:

```
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 and then press **Enter**. If you select **Remote**, you are prompted to provide details about the remote database host.

- Enter the fully qualified domain name and password for your Data Warehouse database host. Press **Enter** to accept the default values in each other field:

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

- Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

```
Reports power users password:
```

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

- For the configuration 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]:
```

Press **Enter** to proceed. The **ovirt-engine** service restarts automatically later in the command.

9. Confirm your installation settings:

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

Next Steps

Access the Reports Portal at <http://demo.redhat.com/ovirt-engine-reports>, replacing *demo.redhat.com* with the fully qualified domain name of the Manager. If during the Manager installation you selected a non-default HTTP port then append *:port* to the URL, replacing *:port* with the port that you chose.

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

4.5. Migrating Data Warehouse and Reports to Separate Machines

Migrate the Data Warehouse service, the Reports service, or both from the Red Hat Enterprise Virtualization Manager to separate machines. Hosting the Data Warehouse service and the Reports service on separate machines reduces the load on each individual machine, and allows each service to avoid potential conflicts caused by sharing CPU and memory with other processes.

Migrate the Data Warehouse service and connect it with the existing **ovirt_engine_history** database, or optionally migrate the **ovirt_engine_history** database to a new database machine before migrating the Data Warehouse service. If the **ovirt_engine_history** database is hosted on the Manager, migrating the database in addition to the Data Warehouse service further reduces the competition for resources on the Manager machine. You can migrate the database to the same machine onto which you will migrate the Data Warehouse service, or to a machine that is separate from both the Manager machine and the new Data Warehouse service machine.

4.5.1. Migrating the Data Warehouse Database to a Separate Machine

Optionally migrate the **ovirt_engine_history** database before you migrate the Data Warehouse service. This procedure uses **pg_dump** to create a database backup, and **psql** to restore the backup on the new database machine. The **pg_dump** command provides flexible options for backing up and restoring databases; for more information on options that may be suitable for your system, see the **pg_dump** manual page.

The following procedure assumes that a PostgreSQL database has already been configured on the new machine. To migrate the Data Warehouse service only, see [Section 4.5.2, “Migrating the Data Warehouse Service to a Separate Machine”](#).



Important

If the existing Data Warehouse database is connected to an existing Reports service, you must reconfigure that service by running **engine-setup** and entering the details of the new Data Warehouse database when prompted. If you do not do this, the Reports service is still connected to the old database, and does not receive any new data.

Procedure 4.9. Migrating the Data Warehouse Database to a Separate Machine

1. On the existing database machine, dump the **ovirt_engine_history** database into a SQL script file:

```
# pg_dump ovirt_engine_history > ovirt_engine_history.sql
```

2. Copy the script file from the existing database machine to the new database machine.
3. Restore the **ovirt_engine_history** database on the new database machine:

```
# psql -d ovirt_engine_history -f ovirt_engine_history.sql
```

The command above assumes that the database on the new machine is also named **ovirt_engine_history**.

4.5.2. Migrating the Data Warehouse Service to a Separate Machine

Migrate a Data Warehouse service that was installed and configured on the Red Hat Enterprise Virtualization Manager to a dedicated host machine. Hosting the Data Warehouse service on a separate machine helps to reduce the load on the Manager machine. Note that this procedure migrates the Data Warehouse service only; to migrate the Data Warehouse database (also known as the **ovirt_engine_history** database) prior to migrating the Data Warehouse service, see [Section 4.5.1, “Migrating the Data Warehouse Database to a Separate Machine”](#).

Installing this scenario involves four key steps:

1. Set up the new Data Warehouse machine.
2. Stop the Data Warehouse service on the Manager machine.
3. Configure the new Data Warehouse machine.
4. Remove the Data Warehouse package from the Manager machine.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager and Data Warehouse on the same machine.
2. To set up the new Data Warehouse machine, you must have the following:
 - ✦ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed.
 - ✦ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlement pools.
 - ✦ The password from the Manager's `/etc/ovirt-engine/engine.conf.d/10-setup-database.conf` file.
 - ✦ Allowed access from the Data Warehouse machine to the Manager database machine's TCP port 5432.
 - ✦ The **ovirt_engine_history** database credentials from the Manager's `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file. If you migrated the **ovirt_engine_history** database using [Section 4.5.1, “Migrating the Data Warehouse Database to a Separate Machine”](#), retrieve the credentials you defined during the database setup on that machine.

Procedure 4.10. Step 1: Setting up the New Data Warehouse Machine

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

```
# subscription-manager register
```

2. Find entitlement pools containing the channels required to install Data Warehouse:

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

3. 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:

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

4. Disable all existing repositories:

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

5. Enable the required channels:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-
rpms
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-server-
rpms
```

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

```
# yum update
```

7. Install the *rhev-dwh-setup* package:

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

Procedure 4.11. Step 2: Stopping the Data Warehouse Service on the Manager Machine

1. Stop the Data Warehouse service:

```
# service ovirt-engine-dwhd stop
```

2. If the **ovirt_engine_history** database, the Manager database, or both are hosted on the Manager machine and were configured by a previous version (Red Hat Enterprise Virtualization 3.4 or prior) that was then upgraded, you must allow the new Data Warehouse machine to access them. Edit the `/var/lib/pgsql/data/postgresql.conf` file and modify the **listen_addresses** line so that it matches the following:

```
listen_addresses = '*'
```

If the line does not exist or has been commented out, add it manually.

If one or both databases are hosted on a remote machine, you must manually grant access by editing the **postgres.conf** file on each machine, and adding the **listen_addresses** line, as above. If both databases are hosted on the Manager machine and were configured during a clean setup of Red Hat Enterprise Virtualization Manager 3.5, access is granted by default.

- Restart the postgresql service:

```
# service postgresql restart
```

Procedure 4.12. Step 3: Configuring the New Data Warehouse Machine

- Run the **engine-setup** command to begin configuration of Data Warehouse on the machine:

```
# engine-setup
```

- Press **Enter** to configure Data Warehouse:

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

- Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

- Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

- Answer the following question about the location of the **ovirt_engine_history** database:

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

Type the alternative option as shown above and then press **Enter**.

- Enter the fully qualified domain name and password for your **ovirt_engine_history** database host. Press **Enter** to accept the default values in each other field:

```
DWH database host []: dwh-db-fqdn
DWH database port [5432]:
```

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

- Enter the fully qualified domain name and password for the Manager database machine. Press **Enter** to accept the default values in each other field:

```
Engine database host []: engine-db-fqdn
Engine database port [5432]:
Engine database secured connection (Yes, No) [No]:
Engine database name [engine]:
Engine database user [engine]:
Engine database password: password
```

- Press **Enter** to create a backup of the existing Data Warehouse database:

```
Would you like to backup the existing database before upgrading it?
(Yes, No) [Yes]:
```

The time and space required for the database backup depends on the size of the database. It may take several hours to complete. If you choose not to back up the database here, and **engine-setup** fails for any reason, you will not be able to restore the database or any of the data within it. The location of the backup file appears at the end of the setup script.

- Confirm that you want to permanently disconnect the existing Data Warehouse service from the Manager:

```
Do you want to permanently disconnect this DWH from the engine?
(Yes, No) [No]:
```

- Confirm your installation settings:

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

Procedure 4.13. Step 4: Removing the Data Warehouse Package from the Manager Machine

- Remove the Data Warehouse package:

```
# yum remove rhvm-dwh
```

This step prevents the Data Warehouse service from attempting to automatically restart after an hour.

- Remove the Data Warehouse files:

```
# rm -rf /etc/ovirt-engine-dwh /var/lib/ovirt-engine-dwh
```

The Data Warehouse service is now hosted on a separate machine from that on which the Manager is hosted.

4.5.3. Migrating the Reports Service to a Separate Machine

Migrate a Reports service that was installed and configured on the Red Hat Enterprise Virtualization Manager to a dedicated host machine. Hosting the Reports service on a separate machine helps to reduce the load on the Manager machine. Note that this procedure migrates the Reports service only. The Reports database (also known as the **ovirt_engine_reports** database) cannot be migrated; you must create a new **ovirt_engine_reports** database when you configure Reports on the new machine. Saved ad hoc reports can be migrated from the Manager machine to the new Reports machine. Migrate the Reports service only after the Manager and Data Warehouse have been configured.

Installing this scenario involves three key steps:

1. Configure the new Reports machine.
2. Migrate any saved reports to the new Reports machine.
3. Remove the Reports service from the Manager machine.

Prerequisites

Ensure that you have completed the following prerequisites:

1. You must have installed and configured the Manager and Reports on the same machine.
2. You must have installed and configured Data Warehouse, either on the Manager machine or on a separate machine.
3. To set up the new Reports machine, you must have the following:
 - ✧ A virtual or physical machine with Red Hat Enterprise Linux 6.6 installed
 - ✧ A subscription to the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlement pools
 - ✧ The password from the Data Warehouse machine's `/etc/ovirt-engine-dwh/ovirt-engine-dwhd.conf.d/10-setup-database.conf` file
 - ✧ Allowed access from the Reports machine to the Manager database machine's TCP port 5432

Procedure 4.14. Step 1: Configuring the New Reports Machine

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

```
# subscription-manager register
```

2. Find entitlement pools containing the channels required to install Reports:

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

3. 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:

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

4. Disable all existing repositories:

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

5. Enable the required channels:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-supplementary-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
# subscription-manager repos --enable=jboss-eap-6-for-rhel-6-server-rpms
```

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

```
# yum update
```

7. Install the *rhev-reports-setup* package:

```
# yum install rhv-reports-setup
```

8. Run the **engine-setup** command to begin configuration of Reports on the machine:

```
# engine-setup
```

9. Press **Enter** to configure Reports:

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

10. Press **Enter** to automatically configure the firewall, or type **No** and press **Enter** to maintain existing settings:

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

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

11. Press **Enter** to accept the automatically detected hostname, or enter an alternative hostname and press **Enter**:

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

12. Enter the fully qualified domain name of the Manager machine, and then press **Enter**:

```
Host fully qualified DNS name of the engine server []:
```

13. Answer the following questions about the **ovirt_engine_reports** database. Press **Enter** to allow setup to create and configure a local database:

```
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]:
```

14. Enter the fully qualified domain name and password for your **ovirt_engine_history** database host. Press **Enter** to accept the default values in each other field:

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

15. Press **Enter** to allow setup to sign the Reports certificate and Apache certificate on the Manager via SSH:

```
Setup will need to do some actions on the remote engine server.
Either automatically, using ssh as root to access it, or you will
be prompted to manually perform each such action.
Please choose one of the following:
1 - Access remote engine server using ssh as root
2 - Perform each action manually, use files to copy content around
(1, 2) [1]:
```

16. Press **Enter** to accept the default SSH port, or enter an alternative port number and then press **Enter**:

```
ssh port on remote engine server [22]:
```

17. Enter the root password for the Manager machine:

```
root password on remote engine server manager-fqdn.com:
```

18. Press **Enter** to allow automatic configuration of SSL on Apache:

```
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]:
```

19. Set a password for the Reports administrative users (**admin** and **superuser**). Note that the reports system maintains its own set of credentials that are separate to those used for the Manager:

```
Reports power users password:
```

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

20. Confirm your installation settings:

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

Procedure 4.15. Step 2: Migrating Saved Reports to the New Reports Machine

1. On the Manager machine, export saved ad hoc reports:

```
# export ADDITIONAL_CONFIG_DIR=/var/lib/ovirt-engine-reports/build-  
conf  
# /usr/share/jasperreports-server-pro/buildomatic/js-export.sh --  
uris /organizations/organization_1/adhoc/aru --output-zip file.zip
```

2. Copy the zip file to the new Reports machine:

```
# scp file.zip reports-machine-fqdn:/
```

3. On the new Reports machine, import saved ad hoc reports:

```
# export ADDITIONAL_CONFIG_DIR=/var/lib/ovirt-engine-reports/build-  
conf  
# /usr/share/jasperreports-server-pro/buildomatic/js-import.sh --  
input-zip file.zip
```

Procedure 4.16. Step 3: Removing the Reports Service from the Manager Machine

1. Stop the Reports service:

```
# service ovirt-engine-reportsd stop
```

2. Remove the Reports package:

```
# yum remove rhvm-reports
```

3. Remove the Reports files:

```
# rm -rf /etc/ovirt-engine-reports /var/lib/ovirt-engine-reports
```

4. Remove the Reports database and user. The default name for both is **ovirt_engine_reports**:

```
# su - postgres  
$ psql  
postgres=# drop database ovirt_engine_reports;  
postgres=# drop user ovirt_engine_reports;
```

**Note**

You can configure more than one working Reports instance, and continue to log in and view reports from an older instance; however, the Manager will directly connect to and have SSO with only the last Reports instance that was configured using **engine-setup**. This means that the Administration Portal includes dashboards from and direct links to only the most recent Reports installation.

Chapter 5. Updating the Red Hat Enterprise Virtualization Environment

This chapter covers both updating your Red Hat Enterprise Virtualization environment between minor releases, and upgrading to the next major version. Always update to the latest minor version of your current Red Hat Enterprise Virtualization Manager version before you upgrade to the next major version.

For interactive upgrade instructions, you can also use the **RHEV Upgrade Helper** available at <https://access.redhat.com/labs/rhevupgradehelper/>. This application asks you to provide information about your upgrade path and your current environment, and presents the relevant steps for upgrade as well as steps to prevent known issues specific to your upgrade scenario.

5.1. Updates 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.

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. A. 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
```

- B. 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

```

5.1.2. Updating the Red Hat Enterprise Virtualization Manager

Updates to the Red Hat Enterprise Virtualization Manager are released via the Content Delivery Network. Before installing an update from the Content Delivery Network, ensure you read the advisory text associated with it and the latest version of the *Red Hat Enterprise Virtualization Manager Release Notes* and *Red Hat Enterprise Virtualization Technical Notes* on the [Customer Portal](#). 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 *rhvm-setup* package:

```
# yum update rhvm-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.

5.1.3. Updating Red Hat Enterprise Virtualization Hypervisors

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. If migration is enabled at cluster level, virtual machines are automatically migrated to another host in the cluster; as a result, it is recommended that Hypervisor updates are performed at a time when the Hypervisor's usage is relatively low.

Ensure that the cluster to which the host belongs has sufficient memory reserve in order for its hosts to perform maintenance. If a cluster lacks sufficient memory, the virtual machine migration operation will hang and then fail. You can reduce the memory usage of this operation by shutting down some or all virtual machines before updating the Hypervisor.

It is recommended that administrators update Red Hat Enterprise Virtualization Hypervisors regularly. Important bug fixes and security updates are included in updates. Hypervisors that are not up to date may be a security risk.



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:

```
# subscription-manager repos --enable=rhel-6-server-rhev-h-rpms
```

3. Ensure that you have the most recent version of the *rhev-hypervisor6* package installed:

```
# yum update rhev-hypervisor6
```

4. From the Administration Portal, click the **Hosts** tab, and then select the Hypervisor that you intend to upgrade.
 - A. If the Hypervisor requires updating, an alert message in the details pane indicates that a new version of the Red Hat Enterprise Virtualization Hypervisor is available.
 - B. If the Hypervisor does not require updating, no alert message is displayed and no further action is required.
5. Click **Maintenance**. If automatic migration is enabled, this causes any virtual machines running on the Hypervisor to be migrated to other hosts. If the Hypervisor is the SPM, this function is moved to another host.
6. Click **Upgrade** to open the **Upgrade Host** confirmation window.
7. Select **rhev-hypervisor.iso**, which is symbolically linked to the most recent Hypervisor image.

- Click **OK** to update and reinstall the Hypervisor. The details of the Hypervisor are updated in the **Hosts** tab, and the status will transition through these stages:

- ✱ **Maintenance**
- ✱ **Installing**
- ✱ **Non Responsive**
- ✱ **Up**

These are all expected, and each stage will take some time.

- Restart the Hypervisor to ensure all updates are correctly applied.

Once successfully updated, the Hypervisor displays a status of **Up**. Any virtual machines that were migrated off the Hypervisor are, at this point, able to be migrated back to it. Repeat the update procedure for each Hypervisor in the Red Hat Enterprise Virtualization environment.



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 **Activate**, and the Hypervisor will change to an **Up** status and be ready for use.

5.1.4. Updating Red Hat Enterprise Linux Virtualization Hosts

Red Hat Enterprise Linux hosts use the **yum** command 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. Updating a host includes stopping and restarting the host. If migration is enabled at cluster level, virtual machines are automatically migrated to another host in the cluster; as a result, it is recommended that host updates are performed at a time when the host's usage is relatively low.

The cluster to which the host belongs must have sufficient memory reserve in order for its hosts to perform maintenance. Moving a host with live virtual machines to maintenance in a cluster that lacks sufficient memory causes any virtual machine migration operations to hang and then fail. You can reduce the memory usage of this operation by shutting down some or all virtual machines before moving the host to maintenance.



Important

Ensure that the cluster contains more than one host before performing an update. Do not attempt to reinstall or update all the hosts at the same time, as one host must remain available to perform Storage Pool Manager (SPM) tasks.

Procedure 5.4. Updating Red Hat Enterprise Linux Hosts

- From the Administration Portal, click the **Hosts** tab and select the host to be updated.
- Click **Maintenance** to place the host into maintenance mode.

3. On the Red Hat Enterprise Linux host machine, run the following command:

```
# yum update
```

4. Restart the host to ensure all updates are correctly applied.

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.

5.2. Upgrading to Red Hat Enterprise Virtualization 3.5

5.2.1. Red Hat Enterprise Virtualization Manager 3.5 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:

- Subscribing to 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.

5.2.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.5

Some of the features provided by Red Hat Enterprise Virtualization 3.5 are only available if your data centers, clusters, and storage have a compatibility version of 3.5.

Table 5.1. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.5

Feature	Description
Paravirtualized random number generator (RNG) device support	This feature adds support for enabling a paravirtualized random number generator in virtual machines. To use this feature, the random number generator source must be set at cluster level to ensure all hosts support and report desired RNG device sources. This feature is supported in Red Hat Enterprise Linux hosts of version 6.6 and higher.

Feature	Description
Serial number policy support	This feature adds support for setting a custom serial number for virtual machines. Serial number policy can be specified at cluster level, or for an individual virtual machine.
Save OVF files on any data domain	This feature adds support for Open Virtualization Format files, including virtual machine templates, to be stored on any domain in a supported pool.
Boot menu support	This feature adds support for enabling a boot device menu in a virtual machine.
Import data storage domains	This feature adds support for users to add existing data storage domains to their environment. The Manager then detects and adds all the virtual machines in that storage domain.
SPICE copy and paste support	This feature adds support for users to enable or disable SPICE clipboard copy and paste.
Storage pool metadata removal	This feature adds support for storage pool metadata to be stored and maintained in the engine database only.
Network custom properties support	This feature adds support for users to define custom properties when a network is provisioned on a host.

5.2.3. Red Hat Enterprise Virtualization 3.5 Upgrade Considerations

The following is a list of key considerations that must be made when planning your upgrade.



Important

Upgrading to version 3.5 can only be performed from version 3.4

To upgrade a previous version of Red Hat Enterprise Virtualization earlier than Red Hat Enterprise Virtualization 3.4 to Red Hat Enterprise Virtualization 3.5, 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 3.3, you must upgrade to the latest minor version of Red Hat Enterprise Virtualization 3.4 before you can upgrade to Red Hat Enterprise Virtualization 3.5.

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 3.5 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.

Red Hat Enterprise Virtualization Manager 3.5 is supported to run on Red Hat Enterprise Linux 6.6

Upgrading to version 3.5 involves also upgrading the base operating system of the machine that hosts the Manager.

5.2.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.5

The following procedure outlines the process for upgrading Red Hat Enterprise Virtualization Manager 3.4 to Red Hat Enterprise Virtualization Manager 3.5. This procedure assumes that the system on which the Manager is installed is subscribed to the entitlements for receiving Red Hat Enterprise Virtualization 3.4 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 repositories required by Red Hat Enterprise Virtualization 3.4 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 5.5. Upgrading to Red Hat Enterprise Virtualization Manager 3.5

1. Subscribe the system on which the Red Hat Enterprise Virtualization Manager is installed to the required entitlements for receiving Red Hat Enterprise Virtualization Manager 3.5 packages:

✧ With RHN Classic:

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.5
```

✧ With Subscription Manager:

```
# subscription-manager repos --enable=rhel-6-server-rhev-3.5-rpms
```

2. Update the *rhev-setup* package:

```
# yum update rhv-setup
```

3. Run the following command and follow the prompts to upgrade the Red Hat Enterprise Virtualization Manager:

```
# engine-setup
```

4. Remove or disable the Red Hat Enterprise Virtualization Manager 3.4 channel to ensure the system does not use any Red Hat Enterprise Virtualization Manager 3.4 packages:

✧ With RHN Classic:

```
# rhn-channel --remove --channel=rhel-x86_64-server-6-rhev-3.4
```

✧ With Subscription Manager:

```
# subscription-manager repos --disable=rhel-6-server-rhev-3.4-rpms
```

5. Update the base operating system:

```
# yum update
```

5.3. Upgrading to Red Hat Enterprise Virtualization 3.4

5.3.1. 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:

- ✧ Subscribing to 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.

5.3.2. 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 5.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.4

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.

Feature	Description
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.

5.3.3. 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.

Reports and the Data Warehouse are now installed via `engine-setup`

From Red Hat Enterprise Virtualization 3.4, the Reports and Data Warehouse features are configured and upgraded using the **`engine-setup`** command. If you have configured the Reports and Data Warehouse features in your Red Hat Enterprise Virtualization 3.3 environment, you must install the *rhev-reports-setup* and *rhev-dwh-setup* packages prior to upgrading to Red Hat Enterprise Virtualization 3.4 to ensure these features are detected by **`engine-setup`**.

5.3.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.4

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 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 repositories 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 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:

```
# subscription-manager repos --enable=rhel-6-server-rhev-3.4-rpms
```

2. Update the *rhev-setup* package:

```
# yum update rhv-setup
```

3. Run the following command and follow the prompts to upgrade the Red Hat Enterprise Virtualization Manager:

```
# engine-setup
```

4. Remove or disable the Red Hat Enterprise Virtualization Manager 3.3 repositories 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:

```
# subscription-manager repos --disable=rhel-6-server-rhev-3.3-rpms
```

5. Update the base operating system:

```
# yum update
```

5.4. Upgrading to Red Hat Enterprise Virtualization 3.3

5.4.1. Red Hat Enterprise Virtualization Manager 3.3 Upgrade Overview

Upgrading Red Hat Enterprise Virtualization Manager is a straightforward process that comprises three main steps:

- ✧ Subscribing to 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.

5.4.2. Features Requiring a Compatibility Upgrade to Red Hat Enterprise Virtualization 3.3

Some of the features in Red Hat Enterprise Virtualization are only available if your data centers, clusters, and storage have a compatibility version of **3.3**.

Table 5.3. 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.

Feature	Description
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.
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.

5.4.3. 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 `rhev-m-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 **rhev-m-upgrade** command used during the installation process has been updated to **engine-setup** and is now obsolete.

5.4.4. Upgrading to Red Hat Enterprise Virtualization Manager 3.3

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 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 repositories 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 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 **subscription-manager** command to enable the repository in your **yum** configuration.

```
# subscription-manager repos --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. 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 base operating system:

```
# 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.

3. Update the *rhev-setup* package to ensure you have the most recent version of **engine-setup**.

```
# yum update rhvm-setup
```

4. 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/articles/408623
Would you like to continue with the upgrade? (Yes, No)
[Yes]:
```

5. Remove Red Hat Enterprise Virtualization Manager 3.2 repositories to ensure the system does not use any Red Hat Enterprise Virtualization Manager 3.2 packages.

Subscription Manager

Use the **subscription-manager** command to disable the Red Hat Enterprise Virtualization 3.2 repository in your **yum** configuration.

```
# subscription-manager repos --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
```

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.

5.5. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

5.5.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.2

Upgrading Red Hat Enterprise Virtualization Manager to version 3.2 is performed using the **rhev-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 **rhev-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. Ensure that the system is subscribed to the required entitlements to receive Red Hat Enterprise Virtualization Manager 3.2 packages. This procedure assumes that the system is already subscribed to required 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-rhev-3.2-rpms** repository associated with the **Red Hat Enterprise Virtualization** entitlement. Use the **subscription-manager** command to enable the repository in your **yum** configuration.

```
# subscription-manager repos --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. 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.

```
# rhn-channel --add --channel=rhel-x86_64-server-6-rhev-3.2
```

2. 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 entitlements.

Certificate-based Red Hat Network

Use the **subscription-manager** command to disable the Red Hat Enterprise Virtualization 3.1 repository in your **yum** configuration. The **subscription-manager** command must be run while logged in as the **root** user.

```
# subscription-manager repos --disable=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 base operating system:

```
# yum update
```

4. To ensure that you have the most recent version of the **rhev-upgrade** command installed you must update the *rhev-setup* package.

```
# yum update rhv-setup
```

5. To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-upgrade** command.

```
# rhv-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)
```

6. 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>.

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.

5.6. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

5.6.1. Upgrading to Red Hat Enterprise Virtualization Manager 3.1

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. Ensure that the system is subscribed to the required 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. 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. 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 **subscription-manager** command to enable the repository in your **yum** configuration. The **subscription-manager** command must be run while logged in as the **root** user.

```
# subscription-manager repos --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.

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. 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 **subscription-manager** command to disable the Red Hat Enterprise Virtualization 3.0 repositories in your **yum** configuration.

```
# subscription-manager repos --disable=rhel-6-server-rhev-3-rpms
```

```
# subscription-manager repos --disable=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 base operating system.

```
# yum update
```

5. To ensure that you have the most recent version of the **rhev-*upgrade*** command installed you must update the *rhev-setup* package.

```
# yum update rhev-setup
```

6. To upgrade Red Hat Enterprise Virtualization Manager run the **rhev-*upgrade*** 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)
```

7. 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>.

8. 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.

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.

5.7. Post-Upgrade Tasks

5.7.1. 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 **C**lusters 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 **E**dit button.
5. Change the **C**ompatibility **V**ersion to the desired value.
6. Click **O**K to open the **C**hange **C**luster **C**ompatibility **V**ersion confirmation window.
7. Click **O**K 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.

5.7.2. 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.

Part III. Installing Hosts

Chapter 6. Introduction to Hosts

6.1. Workflow Progress - Installing Virtualization Hosts



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.

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 on 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*. See [Red Hat Enterprise Linux 6 Installation Guide](#) or [Red Hat Enterprise Linux 7 Installation Guide](#).
- 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.

7.2. Obtaining the Red Hat Enterprise Virtualization Hypervisor Disk Image

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. Obtaining the Red Hat Enterprise Virtualization Hypervisor 6 Disk Image and Tools



Note

For Red Hat Enterprise Virtualization Hypervisor 7 installation instructions, see <https://access.redhat.com/articles/1168703>

1. Enable the **Red Hat Enterprise Virtualization Hypervisor (v.6 x86_64)** repository. With Subscription Manager, attach a **Red Hat Enterprise Virtualization** entitlement and run the following command:

```
# subscription-manager repos --enable=rhel-6-server-rhev-h-rpms
```

2. Install the *rhev-hypervisor6* package. Alternatively, you can download the ISO file from the [Customer Portal](#).

```
# yum install rhev-hypervisor6
```

3. Install the *livecd-tools* package:

```
# yum install livecd-tools
```



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 link also targets the most recently installed version of the Red Hat Enterprise Virtualization ISO image.

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.

7.3.2. Preparing USB Installation Media Using *livecd-iso-to-disk*

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.

Procedure 7.2. Preparing USB Installation Media Using *livecd-iso-to-disk*

1. Ensure you have the latest version of the Red Hat Enterprise Virtualization Hypervisor disk image:

```
# yum update rhev-hypervisor6
```

2. Write the disk image to a USB storage device.

```
# livecd-iso-to-disk --format --reset-mbr /usr/share/rhev-hypervisor/rhev-hypervisor.iso /dev/sdc
```

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.

7.3.4. Preparing USB Installation Media Using dd on Linux Systems

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 7.1. 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.

7.3.5. Preparing USB Installation Media Using dd on Windows Systems

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 Content Delivery 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 7.2. 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 7.3. 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
   8       16    524288  sdb
```

8. 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 7.4. 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
```

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

1. 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 7.5. 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.

7.4. Installing the Red Hat Enterprise Virtualization Hypervisor**7.4.1. Booting the Hypervisor from USB Installation Media**

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. Reboot the system and the Hypervisor boot screen will be displayed.

7.4.2. Booting the Hypervisor from Optical Installation Media

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 is 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 and the Hypervisor boot screen will be displayed.

7.4.3. Installing the Hypervisor



Important

This procedure details the installation instructions for Red Hat Enterprise Virtualization Hypervisor 6. To install Red Hat Enterprise Virtualization Hypervisor 7, see <https://access.redhat.com/articles/1168703>.

Procedure 7.8. Installing the Hypervisor Interactively

1. Insert the installation media that contains the Hypervisor boot image, and start the machine on which you will install the Hypervisor.
2. From the boot splash screen, press any key to open the boot menu.

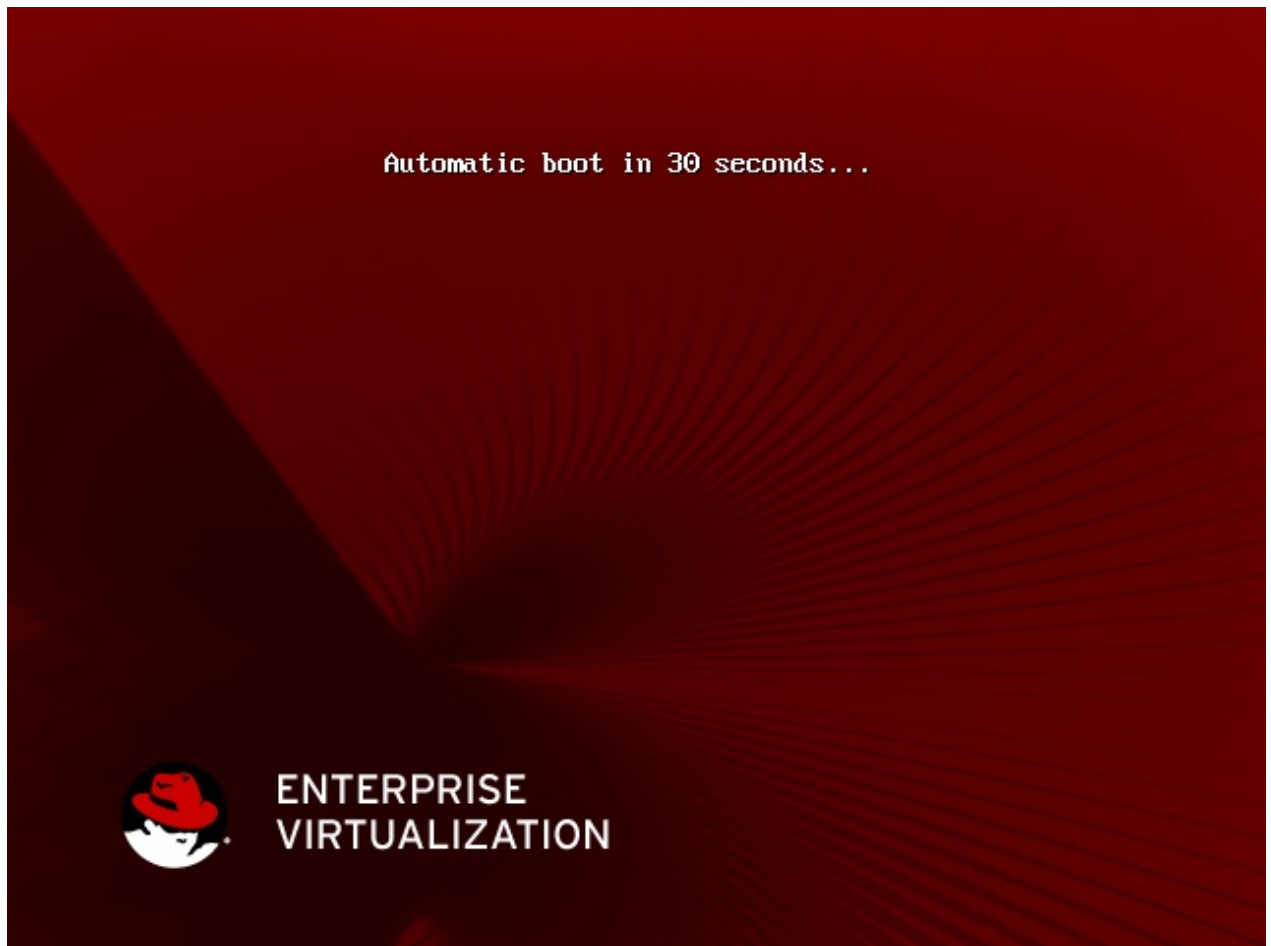


Figure 7.1. The boot splash screen

3. From the boot menu, use the directional keys to select **Install (Basic Video)**.

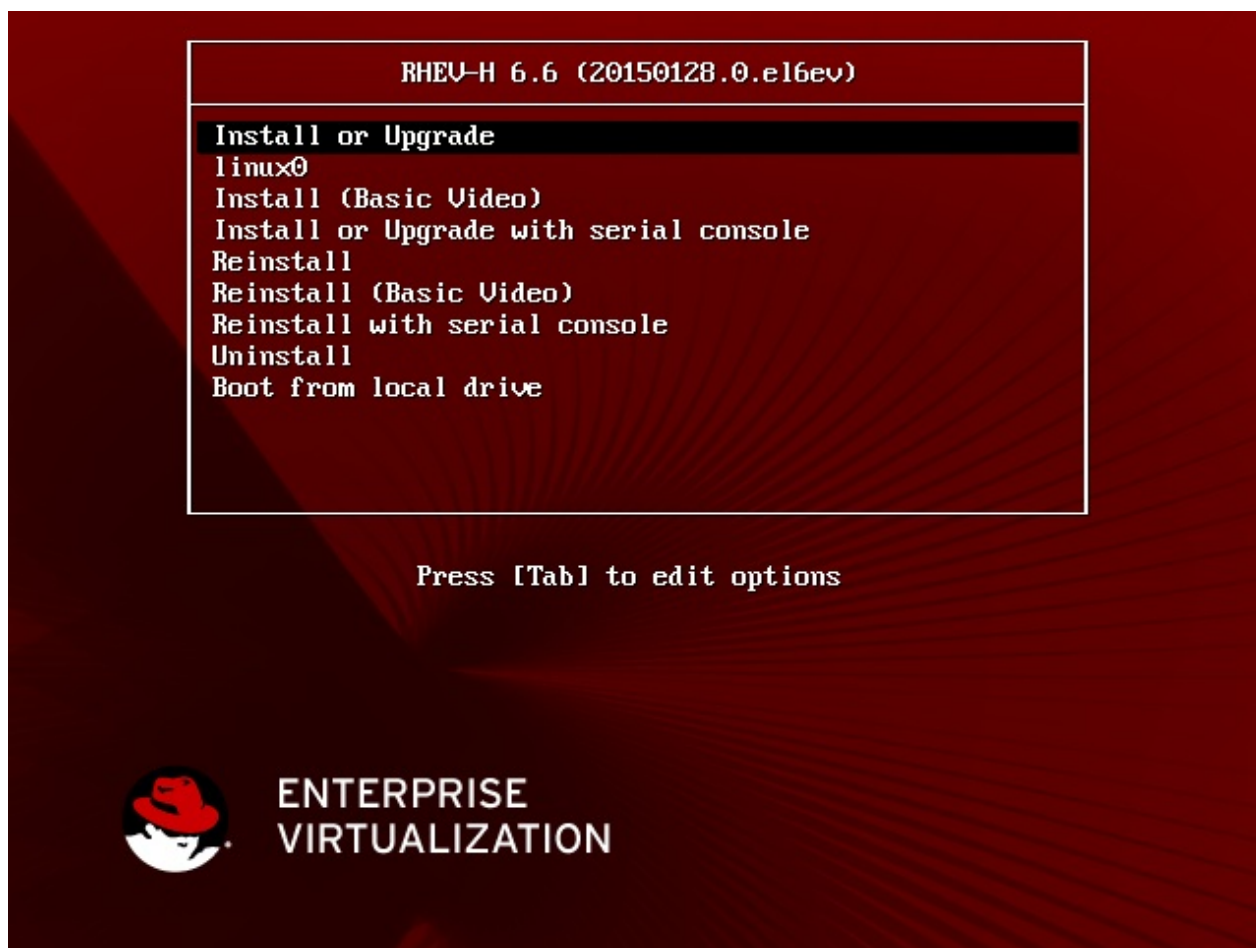


Figure 7.2. The boot menu

4. Customize the keyboard layout to a specific language or location. Use the directional keys to highlight the preferred option, and press **Enter**.

Example 7.6. 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>
```

5. Select the disk from which the Hypervisor will boot. The Hypervisor's boot loader will be installed to the master boot record of the disk that is selected on this screen.

**Important**

The selected disk must be identified as a boot device, and must appear in the boot order either in the system's BIOS or in a pre-existing boot loader.

- A. Select an automatically detected device.
 - a. Select the entry for the disk the Hypervisor is to boot from and press **Enter**.
 - b. Select **<Continue>** and press **Enter**.
- B. Manually select a device.
 - a. Select **Other device** and press **Enter**.
 - b. When prompted to **Please enter the disk to use for booting RHEV-H**, enter the name of the block device from which the Hypervisor should boot.

Example 7.7. Other Device Selection

```
Please enter the disk to use for booting RHEV-H
/dev/sda
```

- c. Press **Enter**.
6. Select the disk or disks on which the Hypervisor itself will be installed. You can select the same device that you selected to act as the boot device if necessary.

**Warning**

All data on the selected storage devices is destroyed.

- a. Select the entry for the disk on which to install the Hypervisor, and press **Enter**. Where other devices are to be used for installation, either solely or in addition to those that are listed automatically, use **Other Device**.
- b. Select **<Continue>**, and press **Enter**.
- c. 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 **<Continue>**, and press **Enter**.

Example 7.8. 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/mmcblk0,/dev/mmcblk1_____
```

7. Configure storage for the Hypervisor.

8. Setting the password and logging in

- a. Select or clear the **Fill disk with Data partition** check box. Clearing this check 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.
 - b. Enter the preferred values for **Swap**, **Config**, and **Logging**.
 - c. If you selected the **Fill disk with Data partition** check box, the **Data** field is automatically set to the value of the **Remaining Space** field. If the check box was cleared, you can enter in the **Data** field a whole number up to the value of the **Remaining Space** field. Entering a value of **-1** fills all remaining space.
8. Set a password to secure local access to the Hypervisor using the administrative user account.
- a. Enter the preferred password in both the **Password** and **Confirm Password** fields.
 - b. Select **<Install>**, and press **Enter**.



Note

If you enter a weak password, the Hypervisor provides a warning; however, you can proceed with the installation using that password. It is recommended that you 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.



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.

7.5. Automated Installation

This section covers the kernel command line parameters for Red Hat Enterprise Virtualization Hypervisors. These parameters can be used to automate installation. The parameters are described in detail and an example parameter string for an automated installation is provided.

This installation method is an alternative to the interactive installation. Using the method covered in this chapter with a PXE server can, with some configuration, deploy multiple Hypervisors without manually accessing the systems.

It is important to understand how the parameters work and what effects they have before attempting automated deployments. These parameters can delete data from existing systems if the system is configured to automatically boot with PXE.

7.5.1. How the Kernel Arguments Work

Below is a description of the Red Hat Enterprise Virtualization Hypervisor start up sequence. This may be useful for debugging issues with automated installation.

1. The **ovirt-early** service sets storage, network and management parameters in the **/etc/default/ovirt** file. These parameters are determined from the kernel arguments passed to the Hypervisor during the boot sequence.
2. The **/etc/init.d/ovirt-firstboot** script determines the type of installation to perform based on the parameters set on the kernel command line or the TUI installation.

7.5.2. Required Parameters

At a minimum, the following parameters are required for an installation:

1. One of the following parameters, depending on the type of installation or reinstallation that you wish to perform:
 - a. **install**, to begin an installation (even if it detects an existing installation).
 - b. **reinstall**, to remove a current installation and begin a completely clean reinstall.
 - c. **upgrade**, to upgrade an existing installation.
2. The **storage_init** parameter, to initialize a local storage device.
3. The **BOOTIF** parameter, to specify the network interface which the Hypervisor uses to connect to the Manager. When using PXE boot, **BOOTIF** may be automatically supplied by **pxelinux**.

These parameters are discussed in further detail in the sections that follow.

If you want to use Red Hat Enterprise Virtualization Hypervisor with Red Hat Enterprise Virtualization Manager, you must also provide at least one of the following parameters:

adminpw

Allows you to log in with administrative privileges to configure Red Hat Enterprise Virtualization Hypervisor.

management_server

Specifies the management server to be used.

rhevm_admin_password

Specifies the password to be used during the process of adding a host in Red Hat Enterprise Virtualization Manager.

7.5.3. Installing to iSCSI Target Root

To configure a Red Hat Enterprise Virtualization Hypervisor host to use iSCSI storage for the Root/HostVG you must provide the following automatic installation parameters along with the required parameters.

iscsi_install

Specifies that iSCSI should be used to boot. This parameter is added to the boot prompt like so:

```
iscsi_install
```

iscsi_init

Defines the device on the target server that should be used for iSCSI. This parameter is added to the boot prompt like so:

```
iscsi_init=device
```

For example, **`iscsi_init=/dev/sdc`**.

iscsi_target_name

Defines the target on the server. This parameter is added to the boot prompt like so:

```
iscsi_target_name=target
```

For example, **`iscsi_target_name=iqn.shared.root`**.

iscsi_server

Defines the iSCSI server and, if required, the port number. This is defined on the boot prompt like so:

```
iscsi_server=server[:port]
```

For example, **`iscsi_server=192.168.1.5:3260`**.

Example 7.9. iSCSI Boot Example

```
BOOTIF=eth0 storage_init=/dev/sda,/dev/sdc \  
iscsi_install \  
iscsi_init=/dev/sdc \  
iscsi_target_name=iqn.shared.root \  
iscsi_server=192.168.1.5:3260
```

Adding this to the boot prompt would specify that:

- ✦ Boot and HostVG are installed to **`/dev/sda`**; and
- ✦ Root is installed to **`/dev/sdc`**.

To use **`/dev/sdc`** as the location for HostVG, just add it to the value of **`storage_init`**.

7.5.4. Storage Parameters

The following parameters configure local storage devices for installing a Hypervisor.

storage_init

The ***storage_init*** parameter is required for an automated installation; it initializes a local storage device.

Hypervisors use one storage device for local installation. There are several methods for defining which disk to initialize and install on.

- ✳ For USB storage devices, use the ***usb*** parameter to select the disk type. For example:

```
storage_init=usb
```

- ✳ For SCSI hard drives, use the ***scsi*** parameter to select the disk type. For example:

```
storage_init=scsi
```

- ✳ For CCISS devices, use the ***cciss*** parameter to select the disk type. For example:

```
storage_init=cciss
```

- ✳ For hard drives on the ATA bus, including SATA hard drives that may also appear on the SCSI bus, use the ***ata*** parameter to select the disk type. For example:

```
storage_init=ata
```

- ✳ Alternatively, the storage device can be specified by using the Linux device name as the ***storage_init*** parameter. Using device names in the format ***/dev/disk/by-id*** is not supported. ***storage_init*** must use the format ***/dev/mapper/disk*** or ***/dev/disk***. In this example the ***/dev/sda*** device is specified:

```
storage_init=/dev/sda
```

When specifying a ***storage_init*** value of ***usb***, ***scsi***, ***ata***, or ***cciss*** you can also append a serial number to explicitly set which device to use. Determine the serial numbers for all disks attached to the system by running the command in the example below:

Example 7.10. Finding udev Serial Numbers

```
$ for d in /dev/sd?; do echo $d `udevadm info -q env -n $d |  
egrep 'ID_BUS=|ID_SERIAL='`; done  
/dev/sda ID_SERIAL=ST9500325AS_6VE867X1
```

When providing both a storage type and the serial number, ensure that the two values are separated by a colon (:). For example:

```
storage_init=cciss:3600508b100104a3953545233304c0003
```



Note

Consistency of device names following a system restart is not guaranteed. Device names are liable to change.

storage_vol

The ***storage_vol*** parameter is used to partition the storage device set by the ***storage_init*** parameter. After ***storage_vol=***, you can specify the size in megabytes of the following partitions: Boot, Swap, Root, Config, Logging, and Data.

The Boot partition is always 50 MB and cannot be reconfigured. The Root partition is always 512 MB and cannot be reconfigured. The remaining partitions are described in more detail below:

Partitions defined by the *storage_vol* parameter

Swap

The swap partition is used for swapping pages of memory that are not frequently accessed to the hard drive. This frees pages of memory in RAM that are in turn used for pages which are accessed more frequently, increasing performance. The default size of the swap partition is calculated based on the amount of RAM installed in the system and over-commit ratio (default is 0.5). Hypervisors must have a swap partition and the swap partition cannot be disabled by setting its size to 0. The minimum size for the swap partition is 8 MB.

To determine the size of the swap partition, see <https://access.redhat.com/knowledge/solutions/15244>.

Use the formula from the Red Hat Knowledgebase solution above and add storage for the over-commit ratio (RAM multiplied by the over-commit ratio).

Recommended swap + (RAM * over-commit) = swap partition size

Leaving the value empty allows the system to set the recommended value for the swap partition.

Config

The config partition stores configuration files for the Hypervisor. The default and minimum size for the configuration partition is 8 MB.

Logging

The logging partition stores all logs for the Hypervisor. The logging partition requires a minimum of 2048 MB storage. However, it is recommended to allocate more storage to the logging partition if resources permit.

Data

The data partition must be large enough to hold core files for KVM. Core files depend on the RAM size for the guests. The data partition must also be large enough to store kernel dump files, also known as kdump. A kdump file is usually the same size as the host's system RAM. The data partition also stores the

Hypervisor ISO file for Hypervisor upgrades.

The data partition requires a minimum of 512 MB storage. The recommended size is at least 1.5 times as large as the RAM on the host system plus an additional 512 MB. It can be configured to take up all remaining space by giving it a size value of **-1**, or disabled by giving it a size value of **0**.

Partitions can be specified in any order. The syntax for specifying each partition is **size, type**. Each partition specified is separated by a colon (:). To specify a 256MB Swap partition, and a 4096MB Logging partition, the whole parameter is **storage_vol=256, Swap: 4096, Logging**.



Note

The old method of specifying partition sizes is still valid. In the old method, the partition sizes must be given in a particular order, as shown here:

```
storage_vol=BOOT:SWAP:ROOT:CONFIG:LOGGING:DATA
```

However, since the Boot and Root partitions cannot be reconfigured, sizes for these partitions can be omitted, like so:

```
storage_vol=:SWAP::CONFIG:LOGGING:DATA
```

If you fail to specify a size, the partition will be created at its default size. To specify a 256MB Swap partition, and a 4096MB Logging partition, the correct syntax is:

```
storage_vol=:256:::4096:
```

The following is the standard format of the **storage_vol** parameter:

```
storage_vol=256, EFI: 256, Root: 4096, Swap
```

iscsi_name

The **iscsi_name** parameter is used to set the iSCSI Initiator Name. The iSCSI Initiator name is expected to take the form of an iSCSI Qualified Name (IQN). This format is defined by RFC 3720, which is available at <http://tools.ietf.org/html/rfc3720>.

The IQN is made up of the following elements, separated by the **.** character:

- The literal string **iqn**
- The date that the naming authority took control of the domain in *yyyy-mm* format
- The reversed domain name - *demo.redhat.com* becomes *com.redhat.demo*
- Optionally, a storage target name as specified by the naming authority - preceded by a colon

Example 7.11. **iscsi_name**

The following illustrates the IQN for an iSCSI initiator attached to the *demo.redhat.com* domain where the domain was established in July 2011.

```
iscsi_name=iqn.2011-07.com.redhat.demo
```

7.5.5. Networking Parameters

Several networking options are available. The following parameters must be appended for the Hypervisor to automatically install:

- ✧ Setting the IP address or DHCP.
- ✧ Setting the hostname if the hostname is not resolved with DHCP.
- ✧ The interface the Red Hat Enterprise Virtualization Manager network is attached to.

The following list contains descriptions and usage examples for both optional and mandatory parameters.

Networking Parameters

BOOTIF

The **BOOTIF** parameter is required for an automated installation.

The **BOOTIF** parameter specifies the network interface which the Hypervisor uses to connect to the Red Hat Enterprise Virtualization Manager.



Important

When using PXE to boot Hypervisors for installation using the **IPAPPEND 2** directive causes **BOOTIF=<MAC>** to be automatically appended to the kernel arguments. If the **IPAPPEND 2** directive is used it is not necessary to use the **BOOTIF** parameter.

The **BOOTIF** parameter takes arguments in one of three forms:

link

Indicates to use the first interface (as enumerated by the kernel) with an active link. This is useful for systems with multiple network interface controllers but only one plugged in.

eth#

Indicates to use the NIC as determined by the kernel driver initialization order (where # is the number of the NIC). To determine the number boot the Hypervisor and select **Shell** from the Hypervisor Configuration Menu. Use **ifconfig | grep eth*** to list the network interfaces attached to the system. There is no guarantee that on the next reboot the network interface controller will have the same **eth#** mapping.

```
BOOTIF=eth0
```

<MAC>

Indicates to use the MAC address explicitly defined inside the brackets.

ip

The ***ip*** parameter sets the IP address for the network interface controller defined by the ***BOOTIF*** parameter. The ***ip*** parameter accepts either an IP address (in the form 0.0.0.0) or ***dhcp***.

```
ip=192.168.1.1
```

```
ip=dhcp
```

ipv6

The ***ipv6*** parameter is an alias for the ***ip*** parameter. It accepts either ***dhcp*** or ***auto***.

netmask

The ***netmask*** parameter sets the subnet mask for the IP address defined with the ***ip*** parameter.

```
netmask=255.255.255.0
```

gateway

The ***gateway*** parameter sets the Internet gateway.

```
gateway=192.168.1.246
```

dns

The ***dns*** parameter sets the address of up to two DNS servers. Each DNS server address must be separated by a comma.

```
dns=192.168.1.243,192.168.1.244
```

hostname

The ***hostname*** parameter sets the hostname. The hostname must be a fully-qualified and resolvable domain name.

```
hostname=rhev1.example.com
```

ntp

The ***ntp*** parameter sets the address of one or more Network Time Protocol servers. Each NTP server address must be separated by a comma.

```
ntp=192.168.2.253,192.168.2.254
```

vlan

The **vlan** parameter sets the VLAN identifier for the network connected to the Red Hat Enterprise Virtualization Manager. This parameter should be set where VLANs are in use.

```
vlan=vlan-id:
```

For example:

```
vlan=36:
```

bond

The **bond** parameter configures a bond. Each interface name must be separated by a comma.

```
BOOTIF=bond01 bond=bond01:nic1,nic2
```

7.5.6. Red Hat Network (RHN) Parameters

These parameters are used to automatically register the hypervisor host with the Red Hat Network (RHN). At a minimum, either the **rh_n_activationkey** or both the **rh_n_username** and **rh_n_password** parameters must be provided. Where registration is to occur against a satellite server, the **rh_n_url** parameter must be provided.

rh_n_type

Sets the RHN entitlement method for this machine. **sam** sets the entitlement method to Certificate-based RHN, which integrates the Customer Portal, content delivery network, and subscription service (subscription management). **classic** sets the entitlement method to RHN Classic, which uses the traditional channel entitlement model (channel access) to provides a global view of content access but does not provide insight into system-level subscription uses. The default value is **sam**.

rh_n_username

The **rh_n_username** parameter sets the username used to connect to RHN.

```
rhn_username=testuser
```

rh_n_password

The **rh_n_password** parameter sets the password used to connect to RHN.

```
rhn_password=testpassword
```

rh_n_activationkey

The **rh_n_activationkey** parameter sets the activation key used to connect to RHN. Activation keys are used to register systems, entitle them to an RHN service level, and subscribe them to specific channels and system groups, all in one action. If both **rh_n_activationkey** and **rh_n_username** are provided, the **rh_n_activationkey** value will be used.

```
rhn_activationkey=7202f3b7d218cf59b764f9f6e9fa281b
```

rh_n_org

This parameter is used only with SAM. Registers the system to SAM in the same way as `--org org_name --activationkey key_value` when combined with the `rhn_activationkey` parameter on the kernel command line.

```
rhn_org=org_name
```

rhn_url

The `rhn_url` parameter sets the URL of the satellite server used to register the host.

```
rhn_url=https://satellite.example.com
```

rhn_ca_cert

The `rhn_ca_cert` parameter sets the URL of the CA certificate used to connect to the satellite server. If it is not provided, the default value is `rhn_url/pub/RHN-ORG-TRUSTED-SSL-CERT`.

```
rhn_ca_cert=https://satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT
```

rhn_profile

The `rhn_profile` parameter sets the name of the profile to be registered with RHN for this host. The default value is the system hostname.

```
rhn_profile=testhost
```

7.5.7. Authentication Parameters

adminpw

The `adminpw` parameter is used to set the password for the `admin` user. The value provided must already be hashed. All hashing schemes supported by the shadow password mechanism are supported. The recommended way to hash a password for use with this parameter is to run the following command:

```
# openssl passwd -1
```

The `openssl` command will prompt for the password to use. A hashed representation of the password will be returned which can be used as the `adminpw` value.

rootpw

The `rootpw` parameter is used to set a temporary `root` password. A password change is forced the first time `root` logs on to the system. The value provided must already be hashed. All hashing schemes supported by the shadow password mechanism are supported. The recommended way to hash a password for use with this parameter is to run the following command:

```
# openssl passwd -1
```

The `openssl` command will prompt for the password to use. A hashed representation of the password will be returned which can be used as the `rootpw` value.



Important

The **root** password is not set by default and is not supported unless enabled at the request of Red Hat support.

rhevm_admin_password

The ***rhevm_admin_password*** parameter sets a **root** password and enables SSH password authentication. The value provided must already be hashed. All hashing schemes supported by the shadow password mechanism are supported. The recommended way to hash a password for use with this parameter is to run the following command:

```
# openssl passwd -1
```

The **openssl** command will prompt for the password to use. A hashed representation of the password will be returned which can be used as the ***rhevm_admin_password*** value.



Important

Setting this parameter has the side-effect of enabling SSH password authentication, which is unsupported unless enabled at the request of Red Hat support. We recommend disabling SSH password authentication after initial configuration is complete.

ssh_pwauth

The ***ssh_pwauth*** parameter is used to select whether or not password authentication is enabled for SSH connections. Possible values are **0** (disabled) and **1** (enabled). The default value is **0**.

```
ssh_pwauth=1
```



Important

SSH password authentication is disabled by default and is not supported unless enabled at the request of Red Hat support.

7.5.8. Other Parameters

firstboot

The ***firstboot*** parameter indicates that the system should be treated as if there is no existing installation.

The ***reinstall*** parameter is a direct alias of the ***firstboot*** parameter, and can be used interchangeably with ***firstboot***.



Warning

Using the ***firstboot*** parameter erases existing data if a disk on the system has a Volume Group named ***HostVG***. Combining the ***firstboot*** parameter with the ***storage_init*** parameter also erases data on any disks specified with ***storage_init***.

install

The ***install*** parameter indicates that the system should be treated as if there is no existing installation. The ***install*** parameter is intended to be used when booting from CD-ROM, DVD, USB, or PXE media.

cim_enabled

Enables the use of Common Information Model (CIM) management infrastructure.

cim_passwd

Configures a password for your Common Information Model (CIM) management infrastructure.

disable_aes_ni

Disables the AES-NI encryption instruction set. Possible values are ***y*** or ***n***.

kdump_nfs

This parameter configures an NFS server for kdump. The syntax for this parameter is ***kdump_nfs=hostname:nfs_share_path***, for example, ***kdump_nfs=nfshost.redhat.com:/path/to/nfs/share***.

local_boot

The ***local_boot*** parameter is an alias for the ***upgrade*** parameter.

local_boot_trigger

Sets a target URL to check and disables PXE when installation completes successfully, so that the system boots from disk on subsequent boots.

netconsole

The ***netconsole*** parameter sets the address of a server to which kernel messages should be logged. The ***netconsole*** parameter takes an IP address or fully qualified domain name and, optionally, a port (the default port is 6666).

```
netconsole=rhev.example.com:6666
```

nfsv4_domain

The ***nfsv4_domain*** parameter specifies a domain to use for NFSv4.

nocheck

The ***nocheck*** parameter will skip the MD5 check of the installation ISO, which might be time consuming if the media is remote or slow.

management_server

The ***management_server*** parameter sets the address of the Red Hat Enterprise Virtualization Manager. The ***management_server*** parameter takes an IP address or fully qualified domain name and, optionally, a port (the default port is 443).

```
management_server=rhev.example.com:443
```

mem_overcommit

The ***mem_overcommit*** parameter specifies the multiplier to use for adding extra swap to support memory over-commit. The default over-commit value is 0.5.

```
mem_overcommit=0.7
```

qemu_pxe

The ***qemu_pxe*** parameter is used to select which network boot loader is used in virtual machines. Possible values are ***gppe*** and ***etherboot***. For compatibility with Red Hat Enterprise Virtualization Hypervisor 5.4-2.1, the default value is ***etherboot***.

```
qemu_pxe=gppe
```

reinstall

The ***reinstall*** parameter indicates that the system should be treated as if there is no existing installation.

The ***firstboot*** parameter is a direct alias of the ***reinstall*** parameter, and can be used interchangeably with ***reinstall***.



Warning

Using the ***reinstall*** parameter erases existing data if a disk on the system has a Volume Group named ***HostVG***. Combining the ***reinstall*** parameter with the ***storage_init*** parameter also erases data on any disks specified with ***storage_init***.

snmp_password

Enables and configures a password for the Simple Network Management Protocol.

syslog

Configures an rsyslog server at the address specified. You can also specify a port. The syntax is ***syslog=hostname[:port]***.

upgrade

The ***upgrade*** parameter will upgrade the existing hypervisor image to the version provided by the boot media. The hypervisor will be automatically upgraded and rebooted once complete. If a hypervisor image is not yet installed, the image will be installed to the device selected with the ***storage_init*** parameter. When performing an upgrade, the previous boot entry is saved as BACKUP in ***grub.conf***. If the reboot following the upgrade procedure fails, the BACKUP boot entry will be automatically selected as the new default.

uninstall

The ***uninstall*** parameter removes an existing Red Hat Enterprise Virtualization installation. The host volume group will be removed and the system rebooted.

7.5.9. An Automated Hypervisor Installation Example

This example uses the kernel command line parameters for an automated Hypervisor installation.



Important

This example may not work accurately on all systems. The parameter descriptions above should be reviewed and the example modified as appropriate for the systems on which deployment is to occur.

The following is a typical example for installing a Hypervisor with the kernel command line parameters.

In this example, the Manager is located at the hostname: **rhevm.example.com**, and the netconsole server is located on the same machine.

```
:linux storage_init=/dev/sda storage_vol::::: local_boot BOOTIF=eth0
management_server=rhevm.example.com netconsole=rhevm.example.com
```



Note

The kernel parameters can be automatically appended to guests booting over a network with PXE. Automatically installing from PXE is not covered by this guide.

7.6. Configuration

7.6.1. Logging in to the Hypervisor

Log in to the Hypervisor console locally to configure the options required to add the Hypervisor to the Manager.

Procedure 7.9. Logging in to the Hypervisor

1. Start the machine on which the Hypervisor is installed.
2. Enter the user name **admin**, and press **Enter**.
3. Enter the password you set during installation, and press **Enter**.

7.6.2. 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.

7.6.3. 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.
- **<Set Console Path>**: Sets a default console device. Enter a path to a valid console device in the **Console path** field.
- **<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.

7.6.4. The Network Screen

7.6.4.1. 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.

7.6.4.2. Configuring the Host Name

Summary

You can change the host name used to identify the Hypervisor.

Procedure 7.10. 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.

7.6.4.3. Configuring Domain Name Servers

Summary

You can specify up to two domain name servers that the Hypervisor will use to resolve network addresses.

Procedure 7.11. 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.

7.6.4.4. 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 7.12. 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.

7.6.4.5. Configuring Network Interfaces

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 7.13. Configuring Network Interfaces

1. Select a network interface 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. Choose to configure either IPv4 or IPv6.
 - A. Configure a dynamic or static IP address for IPv4:
 - a. Select **DHCP** under **IPv4 Settings** and press the space bar to configure a dynamic IP address.
 - b. Select **Static** under **IPv4 Settings**, press the space bar, and input the **IP Address**, **Netmask**, and **Gateway** that the Hypervisor will use to configure a static IP address.

Example 7.12. 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_____
```

- B. Configure a stateless, dynamic, or static IP for IPv6
 - a. Select **Auto** under **IPv6 Settings** and press the space bar to configure stateless auto configuration.
 - b. Select **DHCP** under **IPv6 Settings** and press the space bar to configure a dynamic IP address.
 - c. Select **Static** under **IPv6 Settings**, press the space bar, and input the **IP Address**, **Prefix Length**, and **Gateway** that the Hypervisor will use to configure a static IP address.

Example 7.13. Static IPv6 Networking Configuration

```
IPv6 Settings
( ) Disabled      ( ) DHCP      (*) Static
IP Address: 2001:db8:1::ab9:C0A8:103_  Prefix Length:
```



```
64_____
Gateway      2001:db8:1::ab9:1_____
```

3. Enter a VLAN identifier in the **VLAN ID** field to configure a VLAN for the device.
4. Select the **<Save>** button and press **Enter** to save the network configuration.

7.6.5. 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 7.14. 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.

7.6.6. The Keyboard Screen

Summary

The **Keyboard** screen allows you to configure the keyboard layout used inside the Hypervisor console.

Procedure 7.15. 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.

7.6.7. 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 7.16. 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.

7.6.8. 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 7.17. 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.

7.6.9. 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 7.18. 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. Select an **Interval** to configure **logrotate** to run **Daily**, **Weekly**, or **Monthly**. The default value is **Daily**.
3. 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**.
4. 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**.
5. Select **<Save>** and press **Enter**.

Result

You have configured logging for the Hypervisor.

7.6.10. 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 7.19. Configuring Kernel Dumps

1. Select an option for storing kernel dumps:
 - A. **Local**
 - a. Select the **Local** option and press the space bar to store kernel dumps on the local system.

B. 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.
- c. Enter an **SSH Key URL (optional)**.

C. 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.

7.6.11. The Remote Storage Screen

Use the **Remote Storage** screen to specify a remote iSCSI initiator or NFS share to use as storage.

Procedure 7.20. 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 7.14. iSCSI Initiator Name

```
iSCSI Initiator Name:
iqn.1994-05.com.redhat:5189835eeb40_____
```

Example 7.15. NFS Path

```
NFSv4 Domain (example.redhat.com):
example.redhat.com_____
```

2. Select **<Save>** and press **Enter**.

7.6.12. 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 -1**: Lists partition layout on all block devices.
- » **lsblk**: Lists information on all block devices.

7.6.13. 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
None	The system is disabled from using any tuned profile.
virtual-host	Based on the enterprise-storage profile, virtual-host decreases the swappiness of virtual memory and enables more aggressive writeback of dirty pages.
virtual-guest	A profile optimized for virtual machines.
throughput-performance	A server profile for typical throughput performance tuning.
spindown-disk	A strong power-saving profile directed at machines with classic hard disks.
server-powersave	A power-saving profile directed at server systems.
latency-performance	A server profile for typical latency performance tuning.
laptop-battery-powersave	A high-impact power-saving profile directed at laptops running on battery.
laptop-ac-powersave	A medium-impact power-saving profile directed at laptops running on AC.
enterprise-storage	A server profile to improve throughput performance for enterprise-sized server configurations.
desktop-powersave	A power-saving profile directed at desktop systems.
default	The default power-saving profile. This is the most basic power-saving profile. It only enables the disk and CPU plug-ins.

7.6.14. 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 7.21. Configuring a Hypervisor Management Server

- ✱ A. 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**.
- B. 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**.

7.6.15. The Hosted Engine Screen

To set up a self-hosted engine on Red Hat Enterprise Virtualization Hypervisors (RHEV-H), use the **Hosted Engine** screen. Self-hosted engine on Red Hat Enterprise Virtualization Hypervisor is currently supported on version 6.7 and 7.1 and above. To set up a self-hosted engine on Red Hat Enterprise Linux hosts, see [Section 3.1, “About the Self-Hosted Engine”](#) for more information.

Prerequisites:

- ✦ You must have prepared either NFS or iSCSI storage for your self-hosted engine environment. The storage share must be at least 60 GB. See [Section 12.3, “Preparing NFS Storage”](#) for more information on preparing NFS storage and setting the appropriate permissions.
- ✦ You must have a fully qualified domain name prepared for your Manager and Hypervisor host. Forward and reverse lookup records must both be set in the DNS.

**Note**

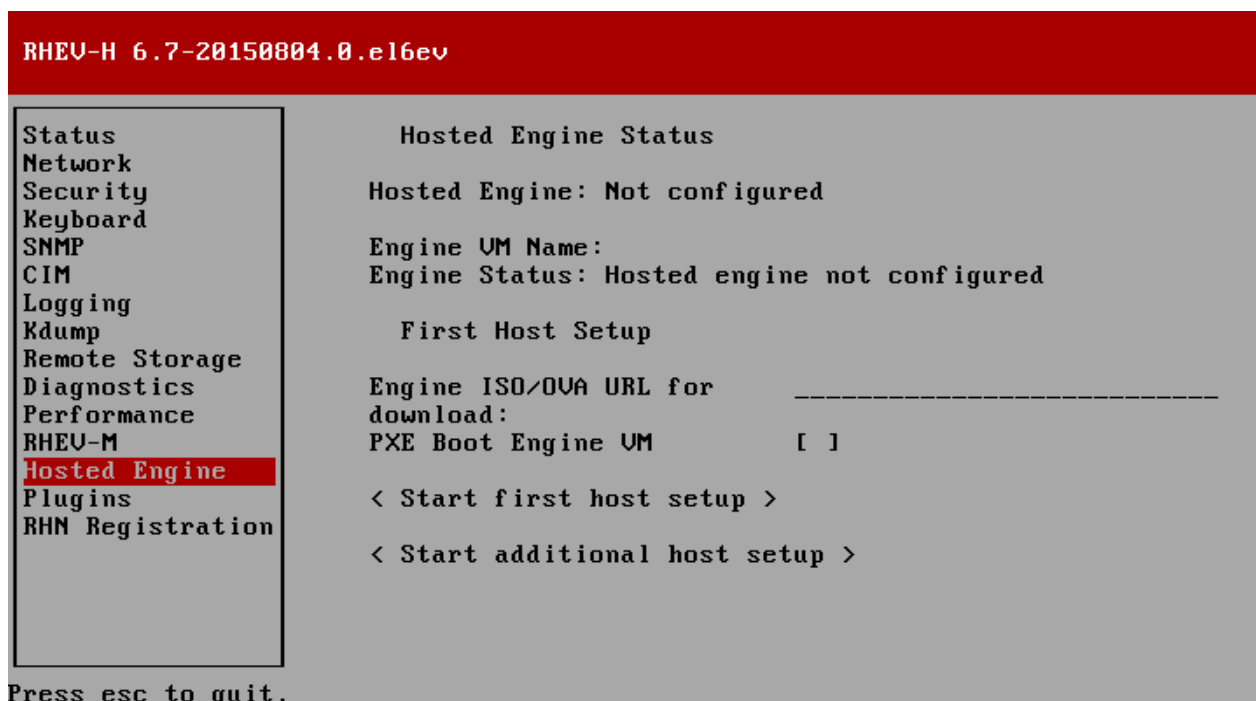
For evaluation purposes, you can use the **/etc/hosts** file for name resolution.

- ✦ The hypervisor has not been previously configured for a Manager.
- ✦ You must have enabled SSH authentication password authentication in the **Security** screen.
- ✦ If you are using the RHEV-M Virtual Appliance for the Manager virtual machine installation and configuration, the **/tmp** directory must be at least 60 GB, and the appliance must be accessible from the hypervisor via HTTP. Download the RHEV-M Virtual Appliance from the [Customer Portal](#).
- ✦ Red Hat Enterprise Virtualization Hypervisor does not support graphical applications, so you must have access to a system that supports graphical applications and has the *virt-viewer* package installed to be able to connect to the Manager virtual machine and complete setup. The *virt-viewer* package is available in standard Red Hat Enterprise Linux repositories.

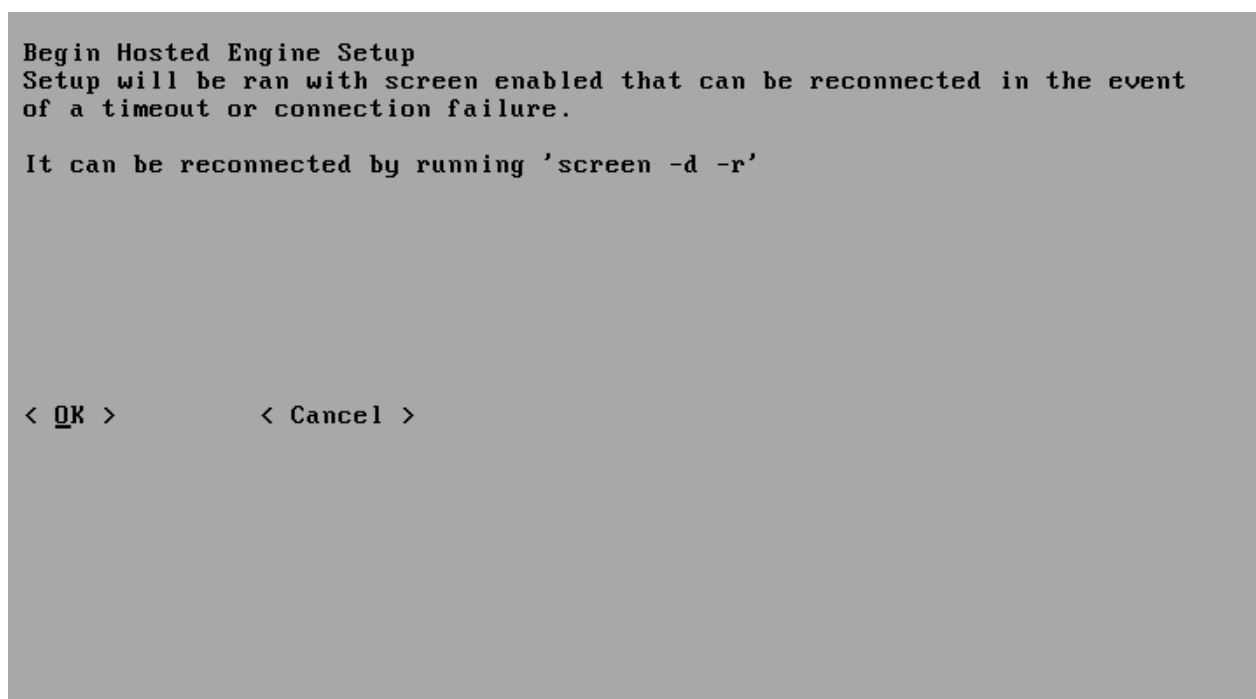
Procedure 7.22. Setting Up Self-Hosted Engine on RHEV-H

This example configures a self-hosted engine on a Red Hat Enterprise Virtualization Hypervisor 6.7 using the RHEV-M Virtual Appliance for Manager virtual machine installation.

1. Specify the Red Hat Enterprise Virtualization Manager virtual machine installation method:
 - A. To install using the RHEV-M Virtual Appliance, enter the URL to the appliance. For example: **http://file.domain.com/rhev-m-appliance.ova**. This is the recommended Manager installation method.
 - B. To install using PXE boot, select the **PXE Boot Engine VM** option.



- Click **<Start first host setup>** and press **Enter**. Click **OK** to start the **hosted-engine** deployment script.



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`



Important

The share must be accessible from the hypervisor and must be owned by user **vdsm** and group **kvm**.

- 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.



Note

Configuring a bonded and vlan-tagged network interface as the management bridge is currently not supported. To work around this issue, see <https://access.redhat.com/solutions/1417783> for more information.

```
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 Virtual Machine

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

```
[ INFO ] Checking OVF archive content
[ INFO ] Checking OVF XML content
```

```
Please specify an alias for the Hosted Engine image [hosted_engine]:
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]:
You may specify a MAC address for the VM or accept a randomly
generated default [00:16:3e:77:b2:a4]:
Please specify the console type you would like to use to connect to
the VM (vnc, spice) [vnc]:
```

6. Configuring the Hosted Engine

Specify a name for the hypervisor 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 the Manager virtual machine.

```
Enter the name which will be used to identify this host inside the
Administrator Portal [hosted_engine_1]:
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]:
```

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.

8. Creating the Manager Virtual Machine

The script creates the Manager virtual machine and provides connection details.

```
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 - engine installation is complete
- (2) Power off and restart the VM
- (3) Abort setup
- (4) Destroy VM and abort setup

```
(1, 2, 3, 4)[1]:
```

9. On a machine that supports graphical applications and has *virt-viewer* installed, connect to the Manager virtual machine. Enter the temporary password.

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

10. Use the RHEV-M Virtual Appliance setup utility to set the root password and change the default authentication and keyboard configuration as necessary. You will not be able to complete RHN registration at this stage as there will be no network connection.



11. Run the following command to complete your Red Hat Enterprise Virtualization Manager setup:

```
# engine-setup --offline --config-append=rhev-m-setup-answers
```

12. Synchronizing the Host and the Manager

Return to the hypervisor 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 ] 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. Shutdown the Manager virtual machine.

```
# shutdown -h now
```

14. Confirm that setup is complete and hit the **Enter** key to return to the Hypervisor console.

```
[ INFO ] Enabling and starting HA services
        Hosted Engine successfully set up
[ INFO ] Stage: Clean up
[ INFO ] Generating answer file '/var/lib/ovirt-hosted-engine-
setup/answers/answers-2015xx.conf'
[ INFO ] Generating answer file '/etc/ovirt-hosted-
engine/answers.conf'
[ INFO ] Stage: Pre-termination
[ INFO ] Stage: Termination

[screen is terminating]
Hit <Return> to return to the TUI
```

After the first host is configured, the Red Hat Enterprise Virtualization Manager name and status are displayed in the **Hosted Engine** screen. This may take a few minutes to appear. Continue to set up another host using the **<Start additional host setup>**.

7.6.16. 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.

7.6.17. 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:

A. 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 7.16. Red Hat Network Configuration

☒ RHN ☐ Satellite ☐ SAM

URL:

—

CA URL:

—

B. 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 7.17. 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

C. 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 7.18. Subscription Asset Manager Configuration

```
( ) RHN      ( ) Satellite    (X) 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 Subscription Asset Manager.

7.7. Adding Hypervisors to Red Hat Enterprise Virtualization Manager

7.7.1. Using the Hypervisor

If the Hypervisor was configured with the address of the Red Hat Enterprise Virtualization Manager, the Hypervisor 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 **RHEV-M** screen during configuration.

Both modes of configuration are supported from the **RHEV-M** screen in the Hypervisor user interface. However, adding the Hypervisor manually is the recommended option.

7.7.2. Approving a Registered Hypervisor

Approve a Hypervisor that has been registered using the details of the Manager.

Procedure 7.24. Approving a Registered Hypervisor

1. From the Administration Portal, click the **Hosts** tab, and then click the host to be approved. The host is currently listed with the status of **Pending Approval**.
2. Click **Approve** to open the **Edit and Approve Hosts** window. You can use the window to specify a name for the Hypervisor, fetch its SSH fingerprint before approving it, and configure power management. For information on power management configuration, refer to [Section 8.9.2, “Host Power Management Settings Explained”](#).
3. Click **OK**. If you have not configured power management, you are prompted to confirm whether to proceed without doing so; click **OK**.

7.7.3. Manually Adding a Hypervisor

Summary

Use this procedure to manually add a Hypervisor that has not been configured with the address of the Manager. You must have both the IP address of the machine upon which the Hypervisor was installed and the password that was set on the **RHEV-M** screen during configuration.

Procedure 7.25. Manually Adding a Hypervisor

1. Log in to the Red Hat Enterprise Virtualization Manager Administration Portal.
2. From the **Hosts** tab, click **New**.
3. In the **New Host** window, enter the **Address** of the Hypervisor, and the root **Password** that was set during configuration. Enter a **Name** for the host, and configure power management, where the host has a supported power management card. For information on power management configuration, refer to [Section 8.9.2, “Host Power Management Settings Explained”](#).



Important

Red Hat recommends to keep Red Hat Enterprise Virtualization Hypervisor 6 and Red Hat Enterprise Virtualization Hypervisor 7 in different clusters.

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**.

7.8. Modifying the Red Hat Enterprise Virtualization Hypervisor ISO

7.8.1. 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.

7.8.2. 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 7.26. 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 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.

7.8.3. 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 the edit-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.

7.8.4. 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.

7.8.4.1. 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 7.27. 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.

7.8.4.2. 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 7.19. 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/rhevhl-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 7.20. 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/rhevhl-latest-6.iso
```

7.8.4.3. 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 7.21. Updating a Single Package in the Red Hat Enterprise Virtualization Hypervisor ISO File

```
# edit-node --nogpgcheck --install vdsm --repo
/etc/yum.repos.d/rhevhl.repo /usr/share/rhev-hypervisor/rhevhl-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 7.22. Updating Multiple Packages in the Red Hat Enterprise Virtualization

```
# edit-node --nogpgcheck --install "vdsm,libvirt" --repo
/etc/yum.repos.d/rhev.repo /usr/share/rhev-hypervisor/rhev-latest-
6.iso
```

7.8.5. Modifying the Default ID of Users and Groups

7.8.5.1. 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 7.23. 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 7.24. Modifying the Default ID of Multiple Users

```
# edit-node --uidmod=user1,60 --uidmod=user2,70
```

7.8.5.2. 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 7.25. 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 7.26. Modifying the Default ID of Multiple Groups

```
# edit-node --gidmod=group1,60 --gidmod=group2,70
```

Chapter 8. Red Hat Enterprise Linux Hosts

8.1. Red Hat Enterprise Linux Hosts

You can use a Red Hat Enterprise Linux 6.6, 6.7, or 7 installation on capable hardware as a host. Red Hat Enterprise Virtualization supports hosts running Red Hat Enterprise Linux 6.6, 6.7, or 7 Server AMD64/Intel 64 version with Intel VT or AMD-V extensions. To use your Red Hat Enterprise Linux machine as a host, you must also attach the **Red Hat Enterprise Linux Server** entitlement and the **Red Hat Enterprise Virtualization** entitlement.

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 process as the host and management system establish a connection.

8.2. Host Compatibility Matrix

The following table outlines the supported hypervisor host versions in each compatibility version of Red Hat Enterprise Virtualization 3.5.

Table 8.1.

Supported RHEL or RHEV-H Version	3.0	3.1	3.2	3.3	3.4	3.5
6.2	✓					
6.3	✓	✓				
6.4	✓	✓	✓			
6.5	✓	✓	✓	✓	✓	
6.6					✓	✓
6.7						✓
7.0						✓
7.1						✓
7.2						✓

8.3. Installing Red Hat Enterprise Linux

You must install Red Hat Enterprise Linux 6.6, 6.7, or 7 Server on a system to use it as a virtualization host in a Red Hat Enterprise Virtualization 3.5 environment.

Procedure 8.1. Installing Red Hat Enterprise Linux

1. Download and Install Red Hat Enterprise Linux

Download and install Red Hat Enterprise Linux 6.6, 6.7, or 7 Server on the target virtualization host. See [Red Hat Enterprise Linux 6 Installation Guide](#) or [Red Hat Enterprise Linux 7 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 *vdsm* 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.

**Important**

By default, SELinux is in enforcing mode upon installation. To verify, run **getenforce**. While it is highly recommended to have SELinux in enforcing mode, it is not required for Red Hat Enterprise Virtualization to host virtual machines. Disabling SELinux eliminates a core security feature of Red Hat Enterprise Linux. Problems also occur when migrating virtual machines between hypervisors that have different SELinux modes. For more information, see [Red Hat Enterprise Linux 6 Virtualization Security Guide](#) or [Red Hat Enterprise Linux 7 Virtualization Security Guide](#).

If you need to live migrate virtual machines from a hypervisor that has SELinux disabled to a hypervisor that has SELinux enabled, see the workaround in <https://access.redhat.com/solutions/1982023>.

2. Ensure Network Connectivity

Following successful installation, 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. A. 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
```

- B. 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**.

- Red Hat Enterprise Linux 6:

```
# service NetworkManager stop
# chkconfig NetworkManager off
```

- Red Hat Enterprise Linux 7:

```
# systemctl stop NetworkManager
# systemctl disable NetworkManager
```

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.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
```

8.4. Subscribing to the Required Entitlements

To be used as a hypervisor host, a Red Hat Enterprise Linux host must be registered and subscribed to a number of entitlements using Subscription Manager. Follow this procedure to register with the Content Delivery Network and attach the **Red Hat Enterprise Linux Server** and **Red Hat Enterprise Virtualization** entitlements to the host.

Procedure 8.2. Subscribing to Required Entitlements using Subscription Manager

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

```
# subscription-manager register
```

2. Find the required subscription pools:

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

3. 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:

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

4. Disable all existing repositories:

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

5. Enable the required repositories:

- A. Red Hat Enterprise Linux 6:

```
# subscription-manager repos --enable=rhel-6-server-rpms
# subscription-manager repos --enable=rhel-6-server-optional-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-mgmt-
agent-rpms
```

- B. Red Hat Enterprise Linux 7:

```
# subscription-manager repos --enable=rhel-7-server-rpms
# subscription-manager repos --enable=rhel-7-server-rhev-mgmt-
agent-rpms
```

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

```
# yum update
```

8.5. Configuring the Virtualization Host Firewall

Summary

Red Hat Enterprise Virtualization requires a number of network ports to 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.

The steps in the following procedure 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.

Procedure 8.3. Configuring the Virtualization Host Firewall

1. Remove existing rules from the firewall 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 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 5900:6923 -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 **5900** to **6923** for guest console connections.
- ✧ Port **16514** for **libvirt** virtual machine migration traffic.
- ✧ Ports **49152** to **49216** for VDSM virtual machine migration traffic.
- ✧ Port **54321** for the Red Hat Enterprise Virtualization Manager.

3. Save the updated firewall configuration

Run the following command to save the updated firewall configuration:

```
# service iptables save
```

4. Enable iptables service

Ensure the **iptables** service is configured to start on boot and has been restarted, or is 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.

8.6. 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 8.4. 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 before exiting.

Result

You have configured **sudo** to allow use by the **root** user.

8.7. 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 8.5. 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**.

A. If **PermitRootLogin** is set to **yes**, or is not set at all, no further action is required.

B. 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.

8.8. 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, with specific entitlements enabled. The physical host must be set up before you can add it to the Red Hat Enterprise Virtualization environment.



Important

Make sure that virtualization is enabled in your host's BIOS settings. For information on changing your host's BIOS settings, refer to your host's hardware documentation.

Procedure 8.6. 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.



Important

Red Hat recommends to keep Red Hat Enterprise Linux 6 and Red Hat Enterprise Linux 7 hosts in different clusters.

4. Enter the **Name**, **Address**, and **SSH Port** of the new host.
5. Select an authentication method to use with the host.
 - A. Enter the root user's password to use password authentication.
 - B. 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 disable use of JSON protocol.



Note

With Red Hat Enterprise Virtualization 3.5, the communication model between the Manager and VDSM now uses JSON protocol, which reduces parsing time. As a result, the communication message format has changed from XML format to JSON format. Web requests have changed from synchronous HTTP requests to asynchronous TCP requests.

- c. 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**.

Result

The new host displays in the list of hosts with a status of **Installing**, and you can view the progress of the installation in the details pane. After installation is complete, the status updates to **Reboot**. The host must be activated for the status to change to **Up**.

8.9. Explanation of Settings and Controls in the New Host and Edit Host Windows

8.9.1. 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 8.2. General settings

Field Name	Description
Data Center	The data center to which the host belongs. Red Hat Enterprise Virtualization Hypervisor hosts cannot be added to Gluster-enabled clusters.
Host Cluster	The cluster to which the host belongs.

Field Name	Description
Use Foreman Hosts Providers	<p>Select or clear this check box to view or hide options for adding hosts provided by Foreman hosts providers. The following options are also available:</p> <p>Discovered Hosts</p> <ul style="list-style-type: none"> ✎ Discovered Hosts - A drop-down list that is populated with the name of Foreman hosts discovered by the engine. ✎ Host Groups - A drop-down list of host groups available. ✎ Compute Resources - A drop-down list of hypervisors to provide compute resources. <p>Provisioned Hosts</p> <ul style="list-style-type: none"> ✎ Providers Hosts - 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 Provider search filter. ✎ Provider search filter - 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.
Name	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.
Comment	A field for adding plain text, human-readable comments regarding the host.
Address	The IP address, or resolvable hostname of the host.
Password	The password of the host's root user. This can only be given when you add the host; it cannot be edited afterwards.
SSH PublicKey	Copy the contents in the text box to the <code>/root/.known_hosts</code> file on the host to use the Manager's ssh key instead of using a password to authenticate with the host.
Automatically configure host firewall	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 Advanced Parameter .
Use JSON protocol	This is enabled by default. This is an Advanced Parameter .

Field Name	Description
SSH Fingerprint	You can fetch the host's SSH fingerprint, and compare it with the fingerprint you expect the host to return, ensuring that they match. This is an Advanced Parameter .

8.9.2. 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 8.3. Power Management Settings

Field Name	Description
Kdump integration	Prevents the host from fencing while performing a kernel crash dump, so that the crash dump is not interrupted. Kdump is available by default on new Red Hat Enterprise Linux 6.6 and 7.1 hosts and Hypervisors. If kdump is available on the host, but its configuration is not valid (the kdump service cannot be started), enabling Kdump integration will cause the host (re)installation to fail. If this is the case, see fence_kdump Advanced Configuration in the <i>Red Hat Enterprise Virtualization Administration Guide</i> .
Primary/ Secondary	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.
Concurrent	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"> ✎ 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. ✎ 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.
Address	The address to access your host's power management device. Either a resolvable hostname or an IP address.

Field Name	Description
User Name	User account with which to access the power management device. You can set up a user on the device, or use the default user.
Password	Password for the user accessing the power management device.
Type	<p>The type of power management device in your host.</p> <p>Choose one of the following:</p> <ul style="list-style-type: none"> ✧ apc - APC MasterSwitch network power switch. Not for use with APC 5.x power switch devices. ✧ apc_snmp - Use with APC 5.x power switch devices. ✧ bladecenter - IBM Bladecenter Remote Supervisor Adapter. ✧ cisco_ucs - Cisco Unified Computing System. ✧ drac5 - Dell Remote Access Controller for Dell computers. ✧ drac7 - Dell Remote Access Controller for Dell computers. ✧ eps - ePowerSwitch 8M+ network power switch. ✧ hpblade - HP BladeSystem. ✧ ilo, ilo2, ilo3, ilo4 - HP Integrated Lights-Out. ✧ ipmilan - Intelligent Platform Management Interface and Sun Integrated Lights Out Management devices. ✧ rsa - IBM Remote Supervisor Adapter. ✧ rsb - Fujitsu-Siemens RSB management interface. ✧ wti - WTI Network Power Switch.
Port	The port number used by the power management device to communicate with the host.
Options	<p>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.</p> <p>For Red Hat Enterprise Linux 7 hosts, if you are using <code>cisco_ucs</code> as the power management device, you also need to append ssl_insecure=1 to the Options field.</p>
Secure	Select 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
Source	Specifies whether the host will search within its cluster or data center for a fencing proxy. Use the Up and Down buttons to change the sequence in which the resources are used.
Disable policy control of power management	Power management is controlled by the Cluster Policy of the host's cluster . 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. Select this check box to disable policy control.

8.9.3. 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 8.4. SPM settings

Field Name	Description
SPM Priority	Defines the likelihood that the host will be given the role of Storage Pool Manager (SPM). The options are Low , Normal , and High 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.

8.9.4. 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 8.5. Console settings

Field Name	Description
Override display address	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).
Display address	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.

Part IV. Basic Setup

Chapter 9. Configuring Data Centers

9.1. Workflow Progress - Planning Your Data Center



9.2. 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.

9.3. 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 creates a **Default** data center at installation. You can also create new data centers that are managed from the same Administration Portal. For example, you may choose to have different data centers for different physical locations, business units, or for reasons of security.

The system administrator, as the superuser, can manage all aspects of the platform (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 automatically has permission to manage all the objects in that data center, including storage domains, clusters, and hosts.

9.4. 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 9.1. 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.

9.5. 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 9.2. 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.

Chapter 10. Configuring Clusters

10.1. 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. The number of hosts and number of virtual machines that belong to a cluster are displayed in the results list under **Host Count** and **VM Count**, respectively.



Note

Red Hat Enterprise Virtualization 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 Gluster Storage Administration Guide*, available at https://access.redhat.com/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*.

10.2. Creating a New Cluster

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 10.1. 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 Type** 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.



Note

For both Intel and AMD CPU types, the listed CPU models are in logical order from the oldest to the newest. If your cluster includes hosts with different CPU models, select the oldest CPU model. For more information on each CPU model, see <https://access.redhat.com/solutions/634853>.

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. Optionally select the **Enable to set VM maintenance reason** check box to enable an optional reason field when a virtual machine is shut down from the Manager, allowing the administrator to provide an explanation for the maintenance.
8. Select either the **/dev/random source** (Linux-provided device) or **/dev/hwrng source** (external hardware device) check box to specify the random number generator device that all hosts in the cluster must use.
9. 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.
10. Click the **Resilience Policy** tab to select the virtual machine migration policy.
11. Click the **Cluster Policy** tab to optionally configure a cluster policy, configure scheduler optimization settings, enable trusted service for hosts in the cluster, enable HA Reservation, and add a custom serial number policy.
12. 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.
13. Click the **Fencing policy** tab to enable or disable fencing in the cluster, and select fencing options.
14. Click **OK** to create the cluster and open the **New Cluster - Guide Me** window.
15. 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.

The new cluster is added to the virtualization environment.

10.3. 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 10.2. 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.

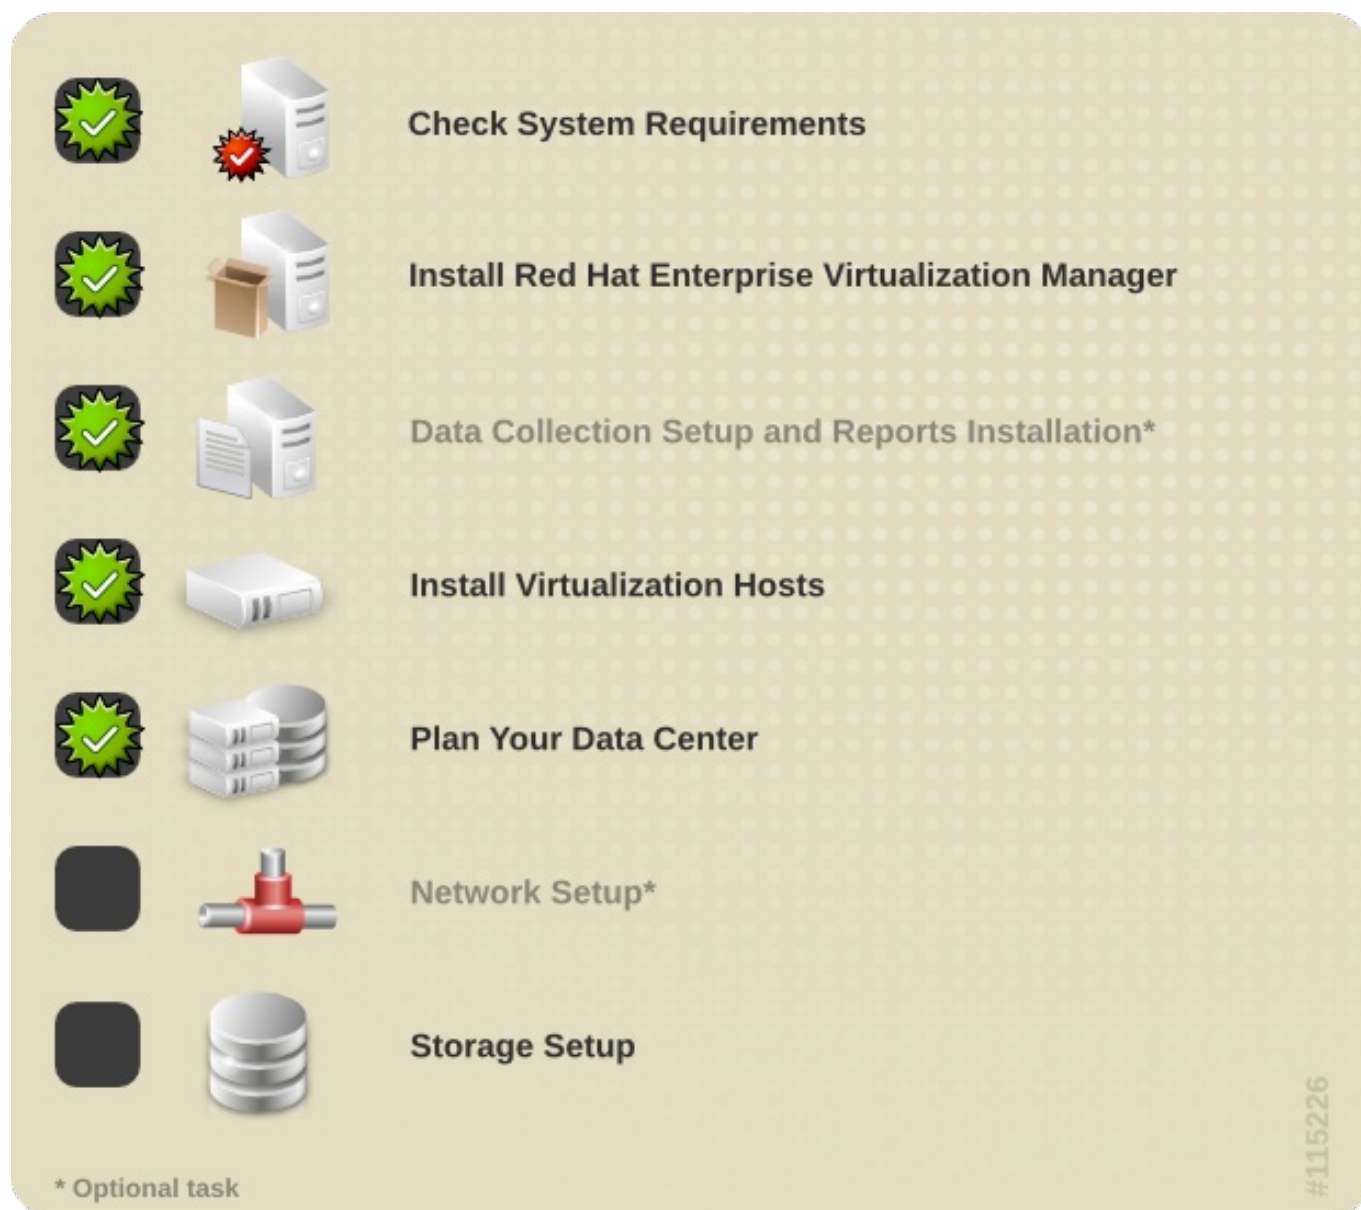


Warning

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.

Chapter 11. Configuring Networking

11.1. Workflow Progress - Network Setup



11.2. 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 **rhevm** network. Incorrect changes to the properties of the **rhevm** 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

11.3. Creating Logical Networks

11.3.1. 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 11.1. 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. A. From the **Data Centers** details pane, click **New** to open the **New Logical Network** window.
B. 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. Optionally select the **Create on external provider** check box. Select the **External Provider** from the drop-down list and provide the IP address of the **Physical Network**.

If **Create on external provider** is selected, the **Network Label**, **VM Network**, and **MTU** options will be removed.
6. Enter a new label or select an existing label for the logical network in the **Network Label** text field.
7. Optionally enable **Enable VLAN tagging**.

8. Optionally disable **VM Network**.
9. Set the **MTU** value to **Default (1500)** or **Custom**.
10. 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.
11. If **Create on external provider** is selected, the **Subnet** tab will be visible. From the **Subnet** tab, enter a **Name**, **CIDR**, and **Gateway** address, and select an **IP Version** for the subnet that the logical network will provide. You can also add DNS servers as required.
12. From the **vNIC Profiles** tab, add vNIC profiles to the logical network as required.
13. 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.

11.4. Editing Logical Networks

11.4.1. 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 11.2. 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.

- Click the **Setup Host Networks** button to open the **Setup Host Networks** window.

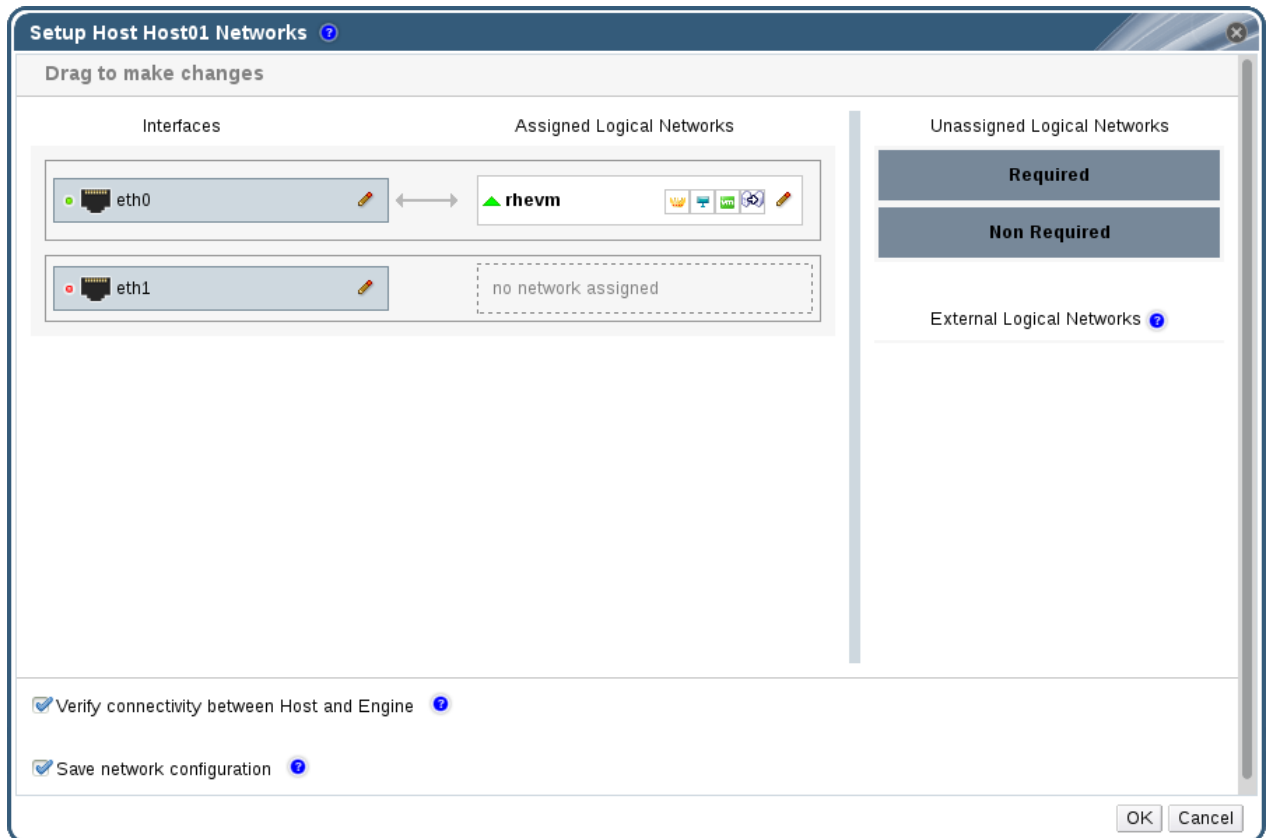


Figure 11.1. The Setup Host Networks window

- 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.

- Configure the logical network:
 - Hover your cursor over an assigned logical network and click the pencil icon to open the **Edit Management Network** window.
 - Select a **Boot Protocol** from:
 - ✦ **None**,
 - ✦ **DHCP**, or
 - ✦ **Static**.

If you selected **Static**, enter the **IP**, **Subnet Mask**, and the **Gateway**.

- To configure a network bridge, click the **Custom Properties** drop-down menu and select **bridge_opts**. Enter a valid key and value with the following syntax: [key]=[value]. Separate multiple entries with a whitespace character. The following keys are valid, with the values provided as examples:

```
forward_delay=1500
```



```

gc_timer=3765
group_addr=1:80:c2:0:0:0
group_fwd_mask=0x0
hash_elasticity=4
hash_max=512
hello_time=200
hello_timer=70
max_age=2000
multicast_last_member_count=2
multicast_last_member_interval=100
multicast_membership_interval=26000
multicast_querier=0
multicast_querier_interval=25500
multicast_query_interval=13000
multicast_query_response_interval=1000
multicast_query_use_ifaddr=0
multicast_router=1
multicast_snooping=1
multicast_startup_query_count=2
multicast_startup_query_interval=3125

```

- d. If your logical network definition is not synchronized with the network configuration on the host, select the **Sync network** check box. A logical network cannot be edited or moved to another interface until it is synchronized.



Note

Networks are not considered synchronized if they have one of the following conditions:

- ✧ The **VM Network** is different from the physical host network.
- ✧ The VLAN identifier is different from the physical host network.
- ✧ A **Custom MTU** is set on the logical network, and is different from the physical host network.

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.

11.4.2. 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 11.1. New Logical Network and Edit Logical Network Settings

Field Name	Description
Name	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.
Description	The description of the logical network. This text field has a 40-character limit.
Comment	A field for adding plain text, human-readable comments regarding the logical network.
Create on external provider	Allows you to create the logical network to an OpenStack Networking instance that has been added to the Manager as an external provider. External Provider - Allows you to select the external provider on which the logical network will be created.
Enable VLAN tagging	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.
VM Network	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.
MTU	Choose either Default , which sets the maximum transmission unit (MTU) to the value given in the parenthesis (), or Custom to set a custom MTU for the logical network. You can use this to match the MTU supported by your new logical network to the MTU supported by the hardware it interfaces with. Enter a numeric value in the text entry field if Custom is selected.
Network Label	Allows you to specify a new label for the network or select from 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.

11.4.3. Editing a Logical Network

Summary

Edit the settings of a logical network.

Procedure 11.3. Editing a Logical Network



Important

A logical network cannot be edited or moved to another interface if it is not synchronized with the network configuration on the host. See [Section 11.4.1, “Editing Host Network Interfaces and Assigning Logical Networks to Hosts”](#) on how to synchronize your 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 **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.

11.4.4. Explanation of Settings in the Manage Networks Window

The table below describes the settings for the **Manage Networks** window.

Table 11.2. Manage Networks Settings

Field	Description/Action
Assign	Assigns the logical network to all hosts in the cluster.
Required	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.

Field	Description/Action
VM Network	A logical network marked "VM Network" carries network traffic relevant to the virtual machine network.
Display Network	A logical network marked "Display Network" carries network traffic relevant to SPICE and to the virtual network controller.
Migration Network	A logical network marked "Migration Network" carries virtual machine and storage migration traffic.

11.4.5. 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 11.4. 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 11.2. 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.

11.4.6. 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 3.5 handles multiple gateways automatically whenever an interface goes up or down.

Procedure 11.5. 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.

11.4.7. 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.

11.5. External Provider Networks

11.5.1. 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 11.6. Importing a Network From an External Provider

1. Click the **Networks** tab.
2. Click the **Import** button to open the **Import Networks** window.

Import Networks

Network Provider: Neutron NP

Provider Networks

<input type="checkbox"/> Name	Provider Network ID
<input type="checkbox"/> network_002	645661fb-c4c2-4925-a401-fe44c828dbbe
<input type="checkbox"/> network_003	534c43bb-e901-4f02-8bbc-2174f0c91723
<input type="checkbox"/> network_004	e7b743f7-ab7a-45ac-bf74-ddb6fa65e192

↓ ↑

Networks to Import

<input type="checkbox"/> Name	Provider Network ID	Data Center	<input checked="" type="checkbox"/> Allow All ?
<input type="checkbox"/> network_001	3a10f9c5-7485-4718-862e-135cfe6545e8	34_DC	<input checked="" type="checkbox"/>

Import Cancel

Figure 11.3. The Import Networks Window

3. 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.
4. 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.
5. It is possible to customize the name of the network that you are importing. To customize the name, click on the network's name in the **Name** column, and change the text.
6. From the **Data Center** drop-down list, select the data center into which the networks will be imported.
7. 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.
8. 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.

11.5.2. 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 Networking instance 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 Enterprise 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.

11.5.3. Configuring Subnets on External Provider Logical Networks

11.5.3.1. 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.

11.5.3.2. Adding Subnets to External Provider Logical Networks

Summary

Create a subnet on a logical network provided by an external provider.

Procedure 11.7. 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.

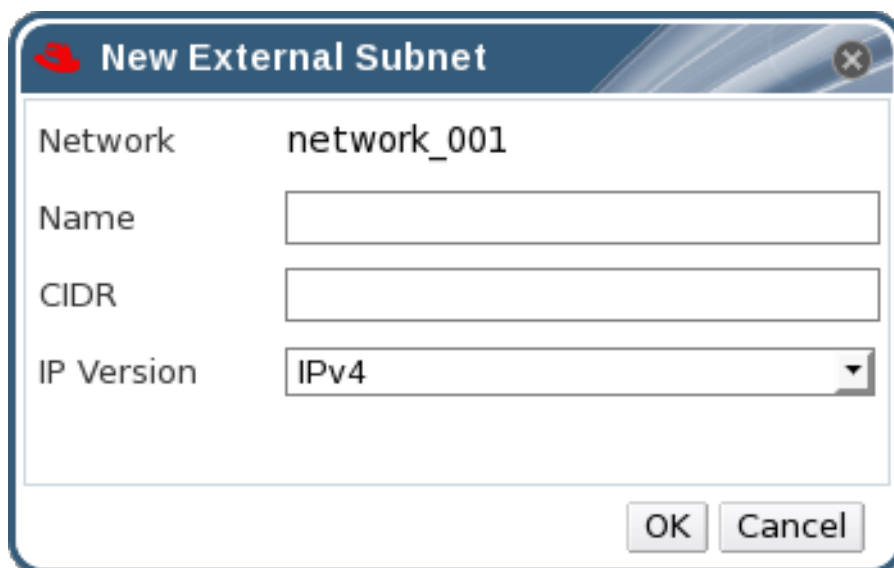


Figure 11.4. The New External Subnet Window

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.

11.5.3.3. Removing Subnets from External Provider Logical Networks

Summary

Remove a subnet from a logical network provided by an external provider.

Procedure 11.8. 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.

11.6. Bonding

11.6.1. 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 11.3. Bonding Scenarios and Their Results

Bonding Scenario	Result
NIC + NIC	<p>The Create New Bond 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 Create New Bond 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>

11.6.2. 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.

11.6.3. 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 11.9. 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.

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.

11.6.4. 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 11.1. 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 11.2. 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 11.3. 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
```

11.7. Removing Logical Networks

11.7.1. Removing a Logical Network

Summary

Remove a logical network from the Manager.

Procedure 11.10. 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.

Chapter 12. Configuring Storage

12.1. Workflow Progress - Storage Setup



12.2. 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.

Red Hat Enterprise Virtualization 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.

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**.

Red Hat Enterprise Virtualization 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. Data domains of multiple types (iSCSI, NFS, FC, POSIX, and Gluster) can be added to the same data center, provided they are all shared, rather than local, domains.

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. ISO domains can only be NFS-based. Only one ISO domain can be added to a data center.
- » **Export Domain:** Export domains are temporary storage repositories that are used to copy and move images between data centers and Red Hat 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. Export domains can only be NFS-based. Only one export domain can be added to a data center.



Important

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).

12.3. Preparing NFS Storage

Set up NFS shares that will serve as a data domain and an export domain on a Red Hat Enterprise Linux 6 server. It is not necessary to create an ISO domain if one was created during the Red Hat Enterprise Virtualization Manager installation procedure.

1. Install `nfs-utils`, the package that provides NFS tools:


```
# yum install nfs-utils
```

2. Configure the boot scripts to make shares available every time the system boots:

```
# chkconfig --add rpcbind
# chkconfig --add nfs
# chkconfig rpcbind on
# chkconfig nfs on
```

3. Start the rpcbind service and the nfs service:

```
# service rpcbind start
# service nfs start
```

4. Create the data directory and the export directory:

```
# mkdir -p /exports/data
# mkdir -p /exports/export
```

5. Add the newly created directories to the **/etc/exports** file. Add the following to **/etc/exports**:

```
/exports/data *(rw)
/exports/export *(rw)
```

6. Export the storage domains:

```
# exportfs -r
```

7. Reload the NFS service:

```
# service nfs reload
```

8. Create the group **kvm**:

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

9. Create the user **vdsm** in the group **kvm**:

```
# useradd vdsm -u 36 -g 36
```

10. Set the ownership of your exported directories to 36:36, which gives vdsm:kvm ownership. This makes it possible for the Manager to store data in the storage domains represented by these exported directories:

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

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

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

```
# chmod 0755 /exports/export
```

12.4. Attaching NFS Storage

Attach an NFS storage domain to the data center in your Red Hat Enterprise Virtualization environment. This storage domain provides storage for virtualized guest images and ISO boot media. This procedure assumes that you have already exported shares. You must create the data domain before creating the export domain. Use the same procedure to create the export domain, selecting **Export / NFS** in the **Domain Function / Storage Type** list.

1. In the Red Hat Enterprise Virtualization Manager Administration Portal, click the **Storage** resource tab.
2. Click **New Domain**.

New Domain ?

Name Description

Data Center Comment

Domain Function / Storage Type Format

Use Host

Export Path

Remote path to NFS export, takes either the form: FQDN:/path or IP:/path e.g. server.example.com:/export/VMs

Advanced Parameters

*** It is recommended to keep the default values in the fields below unchanged.**

☐ Override Default Options

NFS Version

Retransmissions (#)

Timeout (deciseconds)

Additional mount options

OK Cancel

Figure 12.1. The New Domain Window

3. Enter a **Name** for the storage domain.
4. Accept the default values for the **Data Center**, **Domain Function / Storage Type**, **Format**, and **Use Host** lists.
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 **OK**.

The new NFS data domain is displayed in the **Storage** tab with a status of **Locked** until the disk is prepared. The data domain is then automatically attached to the data center.

12.5. Changing the Permissions for the Local ISO Domain

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

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

Procedure 12.1. Changing the Permissions for the Local ISO Domain

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

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

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

3. Apply the changes:

```
# exportfs -ra
```

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

12.6. Attaching the Local ISO Domain to a Data Center

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

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

Procedure 12.2. Attaching the Local ISO Domain to a Data Center

1. In the Administration Portal, click the **Data Centers** resource tab and select the appropriate data center.

2. Select the **Storage** tab in the details pane to list the storage domains already attached to the data center.
3. Click **Attach ISO** to open the **Attach ISO Library** window.
4. Click the radio button for the local ISO domain.
5. Click **OK**.

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

12.7. 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 12.3. 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.

Figure 12.2. 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

- a. 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.

- b. Enter the fully qualified domain name or IP address of the iSCSI host in the **Address** field.
- c. Enter the port to connect to the host on when browsing for targets in the **Port** field. The default is **3260**.
- d. 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**.
- e. Click the **Discover** button.
- f. 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.

12.8. Adding FCP Storage

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, see the [Storage Administration Guide](#) and [DM Multipath Guide](#).

Procedure 12.4. 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

FC-DATA

Description

Data Center

Default (V3)

Comment

Domain Function / Storage Type

Data / Fibre Channel

Format

V3

Use Host

Host_001

<input type="checkbox"/>	LUN ID	Dev. Size	#path	Vendor ID	Product ID	Serial

OK

Cancel

Figure 12.3. 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 through 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. All hosts must have access to the storage device before the storage domain can be configured.

7. The **New Domain** window automatically displays known targets with unused LUNs when **Data / Fibre Channel** is selected as the storage type. Select the **LUN ID** check box to select all of the available LUNs.
8. Click **OK** to create the storage domain and close the window.

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

12.9. Preparing Local Storage

Summary

A local storage domain can be set up on a host. When you set up a 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 12.5. 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 **vdsm** 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.

12.10. 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 12.6. 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.

Figure 12.4. 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 **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.

12.11. 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.

12.12. 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 12.7. 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.

The screenshot shows the 'New Domain' dialog box with the following fields and values:

- Name:** POSIX-DATA
- Description:** (empty)
- Data Center:** Default (V3)
- Comment:** (empty)
- Domain Function / Storage Type:** Data / POSIX compliant FS
- Format:** V3
- Use Host:** (empty)
- Path:** domain.example.com:/Images/POSIX-Data
Path to device to mount / remote export
- VFS Type:** (empty)
- Mount Options:** (empty)

At the bottom right, there are 'OK' and 'Cancel' buttons.

Figure 12.5. 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.

12.13. Enabling Gluster Processes on Red Hat Gluster Storage Nodes

1. In the Navigation Pane, select the **C**lusters 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.

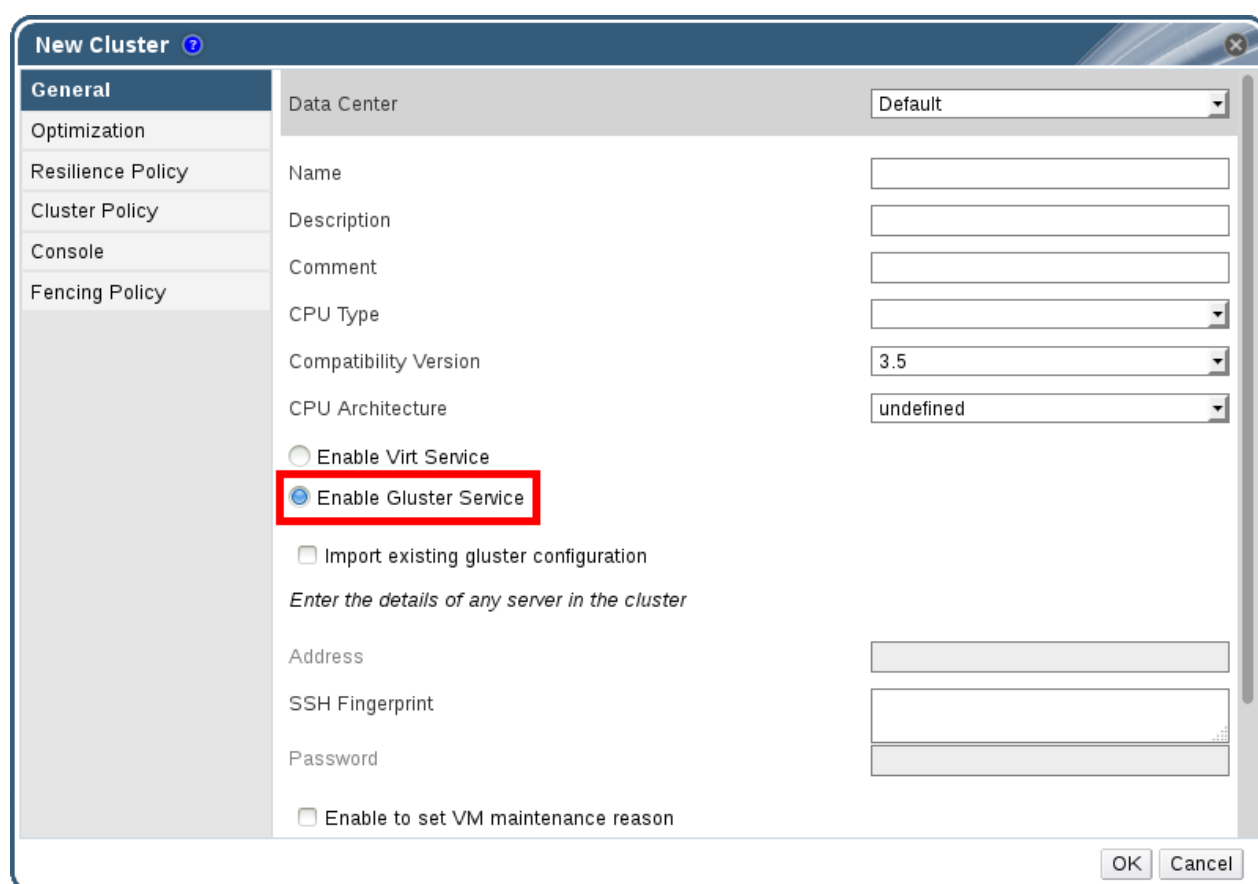


Figure 12.6. Selecting the "Enable Gluster Service" Radio Button

4. Click **OK**.

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

12.14. 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 12.8. 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 12.1. 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 to which the storage domain is attached.

12.15. 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.

12.16. 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 12.2. 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
```

Chapter 13. Configuring Logs

13.1. Red Hat Enterprise Virtualization Manager Installation Log Files

Table 13.1. Installation

Log File	Description
<code>/var/log/ovirt-engine/engine-cleanup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the engine-cleanup 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 engine-setup command detailing the creation and configuration of the rhevm database.
<code>/var/log/ovirt-engine/rhevm-dwh-setup-yyyy_mm_dd_hh_mm_ss.log</code>	Log from the rhevm-dwh-setup command. This is the command used to create the ovirt_engine_history 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 rhevm-reports-setup 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 engine-setup 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.

13.2. Red Hat Enterprise Virtualization Manager Log Files

Table 13.2. Service Activity

Log File	Description
<code>/var/log/ovirt-engine/engine.log</code>	Reflects all Red Hat Enterprise Virtualization Manager GUI crashes, Active Directory lookups, Database issues, and other events.
<code>/var/log/ovirt-engine/host-deploy</code>	Log files from hosts deployed from the Red Hat Enterprise Virtualization Manager.
<code>/var/lib/ovirt-engine/setup-history.txt</code>	Tracks the installation and upgrade of packages associated with the Red Hat Enterprise Virtualization Manager.

13.3. Red Hat Enterprise Virtualization Host Log Files

Table 13.3.

Log File	Description
<code>/var/log/vdsm/libvirt.log</code>	Log file for libvirt .
<code>/var/log/vdsm/spm-lock.log</code>	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.
<code>/var/log/vdsm/vdsm.log</code>	Log file for VDSM, the Manager's agent on the virtualization host(s).
<code>/tmp/ovirt-host-deploy-@DATE@.log</code>	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.

13.4. 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 13.1. 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 the following 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.

13.5. 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 13.2. 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. Select an **Interval** to configure **logrotate** to run **Daily**, **Weekly**, or **Monthly**. The default value is **Daily**.
3. 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**.
4. 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**.
5. Select **<Save>** and press **Enter**.

Result

You have configured logging for the Hypervisor.

Part V. Advanced Setup

Chapter 14. Proxies

14.1. SPICE Proxy

14.1.1. SPICE Proxy Overview

The SPICE Proxy is a tool used to connect SPICE Clients to virtual machines 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 to which the key **SpiceProxyDefault** has been set.



Important

The SPICE Proxy can only be used in conjunction with the standalone SPICE client, and cannot be used to connect to virtual machines using SPICE HTML5 or noVNC.

14.1.2. 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 14.1. 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.

14.1.3. Turning on SPICE Proxy

Summary

This procedure explains how to activate (or turn on) the SPICE proxy.

Procedure 14.2. 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.

14.1.4. Turning Off a SPICE Proxy

Summary

This procedure explains how to turn off (deactivate) a SPICE proxy.

Procedure 14.3. 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.

14.2. Squid Proxy

14.2.1. Installing and Configuring a Squid Proxy

Summary

This section explains how to install and configure a Squid proxy to the User Portal. A Squid proxy server is used as a content accelerator. It caches frequently-viewed content, reducing bandwidth and improving response times.

Procedure 14.4. Configuring a Squid Proxy

1. 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. For this procedure, we assume that they are named **proxy.key** and **proxy.cer**.



Note

The keypair and certificate can also be generated using the certificate authority of the engine. If you already have the private key and certificate for the proxy and do not want to generate it with the engine certificate authority, skip to the next step.

2. Choose a host name for the proxy. Then, choose the other components of the distinguished name of the certificate for the proxy.

**Note**

It is good practice to use the same country and same organization name used by the engine itself. Find this information by logging in to the machine where the Manager is installed and running the following command:

```
# openssl x509 -in /etc/pki/ovirt-engine/ca.pem -noout -subject
```

This command outputs 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
```

3. Log in to the proxy machine and generate a certificate signing request:

```
# openssl req -newkey rsa:2048 -subj '/C=US/O=Example Inc./CN=proxy.example.com' -nodes -keyout proxy.key -out proxy.req
```

**Important**

You must include the quotes around the distinguished name for the certificate. The **-nodes** option ensures that the private key is not encrypted; this means that you do not need to enter the password to start the proxy server.

The command generates 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.

4. To generate the signed certificate, copy the certificate signing request file from the proxy machine to the Manager machine:

```
# scp proxy.req engine.example.com:/etc/pki/ovirt-engine/requests/.
```

5. Log in to the Manager machine and sign the certificate:

```
# /usr/share/ovirt-engine/bin/pki-enroll-request.sh --name=proxy --days=3650 --subject='/C=US/O=Example Inc./CN=proxy.example.com'
```

This signs the certificate and makes it valid for 10 years (3650 days). Set the certificate to expire earlier, if you prefer.

6. The generated certificate file is available in the directory **/etc/pki/ovirt-engine/certs** and should be named **proxy.cer**. On the proxy machine, copy this file from the Manager machine to your current directory:

```
# scp engine.example.com:/etc/pki/ovirt-engine/certs/proxy.cer .
```

7. Ensure both **proxy.key** and **proxy.cer** are present on the proxy machine:

```
# ls -l proxy.key proxy.cer
```

8. Install the Squid proxy server package on the proxy machine:

```
# yum install squid
```

9. Move the private key and signed certificate to a place where the proxy can access them, for example to the **/etc/squid** directory:

```
# cp proxy.key proxy.cer /etc/squid/.
```

10. Set permissions so that the **squid** user can read these files:

```
# chgrp squid /etc/squid/proxy.*
# chmod 640 /etc/squid/proxy.*
```

11. The Squid proxy must verify the certificate used by the engine. Copy the Manager certificate to the proxy machine. This example uses the file path **/etc/squid**:

```
# scp engine.example.com:/etc/pki/ovirt-engine/ca.pem /etc/squid/.
```



Note

The default CA certificate is located in **/etc/pki/ovirt-engine/ca.pem** on the Manager machine.

12. Set permissions so that the **squid** user can read the certificate file:

```
# chgrp squid /etc/squid/ca.pem
# chmod 640 /etc/squid/ca.pem
```

13. If SELinux is in enforcing mode, change the context of port 443 using the **semanage** tool to permit Squid to use port 443:

```
# yum install polycoreutils-python
# semanage port -m -p tcp -t http_cache_port_t 443
```

14. 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
```

15. Restart the Squid proxy server:

```
# service squid restart
```

16. 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 does not support rewriting these headers.



Note

Squid Proxy in the default configuration will terminate its connection after 15 idle minutes. To increase the amount of time before Squid Proxy terminates its idle connection, adjust the **`read_timeout`** option in **`squid.conf`** (for instance **`read_timeout 10 hours`**).

Appendix A. Red Hat Enterprise Virtualization Installation Options

A.1. Configuring a Local Repository for Offline Red Hat Enterprise Virtualization Manager Installation

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

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



Important

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

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

```
# subscription-manager register
```

3. Subscribe the system to all required entitlements:

- a. Find subscription pools containing the repositories 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"
```

- b. 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:

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

- c. Disable all existing repositories:

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

- d. Enable the required repositories:

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

```
# subscription-manager repos --enable=rhel-6-server-
supplementary-rpms
# subscription-manager repos --enable=rhel-6-server-rhev-
3.5-rpms
# subscription-manager repos --enable=jb-eap-6-for-rhel-6-
server-rpms
```

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

```
# yum update
```

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

- a. Install the *vsftpd* package:

```
# yum install vsftpd
```

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

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

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

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

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

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

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

6. Install the *createrepo* package:

```
# yum install createrepo
```

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

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

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

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

```
#!/bin/sh

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

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

9. Install the Manager packages on the offline system. See [Section 2.4.1, “Installing the Red Hat Enterprise Virtualization Manager Packages”](#) for instructions. Packages are installed from the local repository, instead of from the Content Delivery Network.
10. Configure the Manager. See [Section 2.4.4, “Configuring the Red Hat Enterprise Virtualization Manager”](#) for initial configuration instructions.
11. Continue with host, storage, and virtual machine configuration.

A.2. Updating the Local Repository for an Offline Red Hat Enterprise Virtualization Manager Installation

If your Red Hat Enterprise Virtualization Manager is hosted on a system that receives packages via FTP from a local repository, you must regularly synchronize the repository to download package updates from the Content Delivery Network, then update or upgrade your Manager system. Updated packages address security issues, fix bugs, and add enhancements.

1. On the system hosting the repository, synchronize the repository to download the most recent version of each available package:

```
# reposync -l --newest-only /var/ftp/pub/rhevrepo
```

This command may download a large number of packages, and take a long time to complete.

2. Ensure that the repository is available on the Manager system, and then update or upgrade the Manager system. See [Section 5.1.2, “Updating the Red Hat Enterprise Virtualization Manager”](#) for information on updating the Manager between minor versions. See [Section 5.2.4, “Upgrading to Red Hat Enterprise Virtualization Manager 3.5”](#) for information on upgrading the Manager to version 3.5.

Appendix B. Revision History

Revision 3.5-75	Wed 27 Jul 2016	Red Hat Enterprise Virtualization Documentation Team
BZ#1359544 - Updated the links to the Storage Administration Guide and DM Multipath Guide.		
Revision 3.5-74	Mon 18 Apr 2016	Red Hat Enterprise Virtualization Documentation Team
BZ#1320923 - Updated the link to the RHEVM Virtual Appliance. BZ#1306517 - Added an overview and diagrams of the self-hosted engine backup and restoration procedure.		
Revision 3.5-73	Wed 23 Mar 2016	Red Hat Enterprise Virtualization Documentation Team
BZ#1317378 - Added a link to a Customer Portal solution that outlines steps for enabling a new Manager virtual machine and removing a dead Manager virtual machine in scenarios where the Manager database is restored successfully, but the Manager virtual machine appears to be Down and cannot be migrated to another self-hosted engine host.		
Revision 3.5-72	Thu 25 Feb 2016	Red Hat Enterprise Virtualization Documentation Team
BZ#1310139 - Updated the storage requirements for deploying the RHEV-M Virtual Appliance with the RHEV-H-based self-hosted engine.		
Revision 3.5-71	Mon 08 Feb 2016	Red Hat Enterprise Virtualization Documentation Team
BZ#1290096 - Updated the 3.4 and 3.5 upgrade procedures. BZ#1278310 - Added link to required subscriptions in the virtual machine configuration procedure in the self-hosted engine section. BZ#1288672 - Added a note warning that configuring a bonded and vlan-tagged network interface as the management bridge is currently not supported. BZ#1151265 - Added new topic 'Removing a Host from a Self-Hosted Engine Environment'. BZ#1300145 - Added an important warning that SELinux should stay enforcing as problems with migrating hosts can occur if it is disabled.		
Revision 3.5-70	Fri 25 Dec 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1275504 - Clarified the topic 'Migrating to a Self-Hosted Environment'.		
Revision 3.5-69	Tue 08 Dec 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1285598 - Added procedures 'Migrating the Data Warehouse Database to a Remote Database Server' and 'Migrating the Reports Database to a Remote Server Database'. BZ#1276122 - Corrected syntax for automatic VLAN assignment for Red Hat Enterprise Virtualization Hypervisor.		

Revision 3.5-68	Tue 01 Dec 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1285598 - Added procedure 'Migrating the Self-Hosted Engine Database to a Remote Server Database'. Updated the RHEV-M Virtual Appliance download link.		
Revision 3.5-67	Tue 24 Nov 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1281642 - Added a step to disable all repositories after subscribing to a pool id. BZ#1276121 - Updated the automated RHEV-H installation networking parameters section. BZ#1229358 - Updated the Host Compatibility Matrix table. BZ#1283819 - Updated the subscription topics to make sure yum update is run after the required repositories are enabled.		
Revision 3.5-66	Mon 26 Oct 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1259804 - Updated browser requirements. BZ#1175406 - Clarified the Hypervisor upgrade procedure. BZ#1273715 - Removed documentation for an unsupported feature. BZ#1247486 - Updated the upgrade instructions. BZ#1249163 - Added the word optional to the SNMP and CIM port descriptions.		
Revision 3.5-65	Thu 24 Sep 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1224935 - Updated the procedure for configuring an offline repository, and added a procedure for updating the offline repository. BZ#1265293 - Updated the OS version requirements for the Manager.		
Revision 3.5-64	Mon 07 Sep 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1259568 - Added if using cisco_ucs as the power management device on Red Hat Enterprise Linux 7 hosts, ssl_insecure=1 needs to be appended to the options field.		
Revision 3.5-63	Tue 25 Aug 2015	Red Hat Enterprise Virtualization Documentation Team
Minor updates for RHEV 3.5.4.		
Revision 3.5-62	Tue 25 Aug 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1241478 - Added content that covers deployment of self-hosted engine on Red Hat Enterprise Virtualization Hypervisor. BZ#1243235 - Added missing steps to upgrade the self-hosted engine hosts from Red Hat Enterprise Virtualization 3.4 to 3.5. BZ#1253749 - Updated the host matrix table and self-hosted engine content to include support for RHEL 6.7.		

Revision 3.5-61	Tue 04 Aug 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1229797 - Updated information on setting host CPU at the cluster level, and updated information on adding FCP storage. BZ#1248273 - Updated virtualization host storage requirements. BZ#1229486 - Updated SPICE support information. BZ#1213305 - Updated the introduction to storage topic.		
Revision 3.5-60	Wed 08 Jul 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1231025 - Updated considerations for upgrading to 3.5. BZ#1240869 - Updated the list of repositories for RHEL 7 hosts.		
Revision 3.5-59	Fri 03 Jul 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1232136 - New section 'Backing up and Restoring a Self-Hosted Environment' added with six new topics.		
Revision 3.5-58	Wed 24 Jun 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1232140 - Added steps for removing the Data Warehouse and Reports files from the previous machine after migrating.		
Revision 3.5-57	Mon 15 Jun 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1227563 - Added download and supporting documentation links to the Hypervisor installation chapter.		
Revision 3.5-56	Tue 02 Jun 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1216289 - Added a link to the RHEV Upgrade Helper lab application. BZ#1190653 - Updated the creating a new cluster topic.		
Revision 3.5-55	Tue 12 May 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1160742 - Added a step to topics on configuring Data Warehouse and Reports on the Manager about moving self-hosted engine to maintenance mode. BZ#1215998 - Removed an invalid command from the procedure for migrating to a self-hosted environment. BZ#1215283 - Fixed an incorrect chkconfig command.		
Revision 3.5-54	Tue 28 Apr 2015	Red Hat Enterprise Virtualization Documentation Team

[BZ#1188416](#) - Updated subscription and installation media information for self-hosted engine.
[BZ#1207030](#) - Updated the host compatibility matrix to include Red Hat Enterprise Linux 7.1.
[BZ#1205009](#) - Updated the self-hosted engine repo list.
[BZ#1204579](#) - Updated the incorrect subscription-manager command syntax.
[BZ#1172379](#) - Revised instructions for Manager database setup.
[BZ#1204585](#) - Added further explanation of SCSI ALUA.
[BZ#1169192](#) - Updated the PCI device section and updated the required repo list for RHEL 6 hosts.
[BZ#1209333](#) - Checked all repo list and updated terminology for RHSM.
[BZ#1206392](#) - Changed all instances of 'Red Hat Storage' to 'Red Hat Gluster Storage'.
[BZ#1193251](#) - Updated the browser and client requirements.
[BZ#1204170](#) - Updated the use of the Red Hat Enterprise Virtualization Manager Virtualization Appliance with the Hosted Engine.

Revision 3.5-53	Thu 19 Mar 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1190655 - Added documentation for the Red Hat Enterprise Virtualization Manager Virtualization Appliance.		

Revision 3.5-52	Fri 13 Mar 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1195448 - Updated the self-hosted engine repo list.		

Revision 3.5-51	Wed 11 Mar 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1191809 - Added the ipv6 and bond_options parameters for the Red Hat Enterprise Virtualization Hypervisor. BZ#1199474 - Updated the procedure for upgrading the self-hosted engine. BZ#1191810 - Added the disable_aes_ni and nfsv4_domain parameters for the Red Hat Enterprise Virtualization Hypervisor. BZ#1191813 - Updated information on data partition sizing for the Red Hat Enterprise Virtualization Hypervisor. BZ#1122912 - Added documentation for the Diagnostics screen of the hypervisor configuration menu. BZ#1122915 - Added documentation for the Performance screen of the hypervisor configuration menu. BZ#1122919 - Added documentation for the Plugins screen of the hypervisor configuration menu. BZ#1192319 - Added RHEL 7 to the host compatibility matrix table. BZ#1193252 - Updated installation instructions. BZ#1190988 - Added SSH Key URL (optional) field to the Kdump Screen for the Red Hat Enterprise Virtualization Hypervisor. BZ#1190992 - Added Enable SCSI DH_ALUA field to the Remote Storage Screen for the Red Hat Enterprise Virtualization Hypervisor.		

Revision 3.5-50	Fri 27 Feb 2015	Red Hat Enterprise Virtualization Documentation Team
------------------------	------------------------	---

[BZ#1195448](#) - Updated the supported versions of hosts for hosted engine deployment and included additional information on using the 'screen' command for deploying the hosted engine over a network.

[BZ#1156009](#) - Added a procedure for migrating the Reports service to a separate machine.

[BZ#1156015](#) - Added procedures for migrating the Data Warehouse service and Data Warehouse database to separate machines.

[BZ#1193686](#) - Added a note on keeping RHEL 6/RHEV-H 6 and RHEL 7/RHEV-H 7 in different clusters.

[BZ#1172331](#) - Updated upgrade procedures.

[BZ#1190993](#) - Added Organization field to the RHN Registration Screen for the Red Hat Enterprise Virtualization Hypervisor.

[BZ#1190983](#) - Added Set Console Path field to the Status Screen for the Red Hat Enterprise Virtualization Hypervisor.

[BZ#1190986](#) - Added Interval field to the Logging Screen for the Red Hat Enterprise Virtualization Hypervisor.

[BZ#1190994](#) - Added a link to the RHEV-H 7 article.

Revision 3.5-49	Thu 19 Feb 2015	Red Hat Enterprise Virtualization Documentation Team
BZ#1182144 - Added a link that directs users to the History Database Size Calculator tool.		
BZ#1192491 - Corrected a command that enables repositories for Red Hat Enterprise Linux 7 hosts.		
BZ#1156009 - Added a procedure for migrating the Reports service to a separate machine.		
BZ#1156015 - Added procedures for migrating the Data Warehouse service and Data Warehouse database to separate machines.		
Revision 3.5-48	Tue 10 Feb 2015	Andrew Burden
BZ#1075540 - Updated the information on the maintenance modes for the self-hosted engine.		
Revision 3.5-47	Mon 9 Feb 2015	Julie Wu
BZ#1172951 - Added content on automated installation for Red Hat Enterprise Virtualization Hypervisors.		
Revision 3.5-46	Fri 6 Feb 2015	Tahlia Richardson
BZ#1184670 - Updated the Hypervisor section to add a recommendation and change "the oVirt Engine screen" to "the RHEV-M screen".		
Revision 3.5-45	Fri 23 Jan 2015	Lucy Bopf
BZ#1169176 - Updated the list of prerequisites for installing the Manager.		
Revision 3.5-44	Mon 19 Jan 2015	Lucy Bopf
BZ#1169176 - Updated content on the permissions for the local ISO domain.		
Revision 3.5-43	Mon 19 Jan 2015	David Ryan
BZ#1153351 - Updated the supported management client configurations.		
Revision 3.5-42	Tue 13 Jan 2015	Lucy Bopf
BZ#1176795 - Moved Data Warehouse and Reports installation content into the Installation Guide.		
Revision 3.5-41	Tue 06 Jan 2015	Lucy Bopf

[BZ#1121878](#) - Reinstated the History and Reports chapter heading, and added a link that directs users to the Data Warehouse and Reports article set.

Revision 3.5-40	Mon 15 Dec 2014	Andrew Burden
Review and multiple corrections to 'hypervisor' when used as a proper noun.		
Revision 3.5-39	Thurs 11 Dec 2014	Tahlia Richardson
BZ#1172299 - Updated the command for saving iptables rules persistently.		
Revision 3.5-38	Mon 08 Dec 2014	Julie Wu
BZ#1170798 - Updated with RHEV-H storage requirements and updated the Hypervisor installation section to point to the RHEV-H 7.0 installation article.		
BZ#1157205 - Updated the supported RHEL version for the Manager.		
BZ#1124129 - Added a note to highlight the implementation of JSON protocol in 3.5.		
Revision 3.5-37	Wed 26 Nov 2014	Tahlia Richardson
BZ#1149970 - Adjusted the description of firewall port 6100.		
Revision 3.5-36	Tues 18 Nov 2014	Julie Wu
BZ#1164726 - Updated the fixed URL for 'Adding Local Storage' and 'Preparing Local Storage'.		
Revision 3.5-35	Tues 18 Nov 2014	Lucy Bopf
BZ#1121878 - Removed outdated chapter on history and reports, as this content has been expanded and moved to the Customer Portal.		
Revision 3.5-34	Wed 12 Nov 2014	Andrew Dahms
BZ#1044852 - Updated the procedure on installing the Red Hat Enterprise Virtualization Hypervisor.		
Revision 3.5-33	Sun 09 Nov 2014	Laura Novich
BZ#1123921 - Added new option for host network bridging.		
Revision 3.5-32	Fri 07 Nov 2014	Tahlia Richardson
BZ#1149970 - Added rows to the Manager firewall table for ports 6100 and 7410.		
BZ#1157934 - Adjusted the layout of all of the firewall tables.		
Revision 3.5-31	Tue 04 Nov 2014	Lucy Bopf
BZ#1155377 - Revised section on installing and configuring a Squid proxy.		
Revision 3.5-30	Mon 03 Nov 2014	Lucy Bopf
BZ#1138480 - Removed information suggesting that the default data center should not be removed.		
Revision 3.5-29	Tue 28 Oct 2014	Tahlia Richardson
BZ#1150148 - Removed the default storage type question from RHEV-M setup.		
BZ#1123951 - Added iSCSI as a storage options for the self-hosted engine, and removed the topic "Limitations of the Self-Hosted Engine".		
Revision 3.5-28	Tue 28 Oct 2014	Julie Wu
BZ#1154537 - Added an important note on upgrading to the latest minor version before upgrading to the next major version.		
Revision 3.5-27	Tue 21 Oct 2014	Tahlia Richardson

[BZ#1125070](#) - Edit for typos and inconsistencies.

Revision 3.5-26	Mon 20 Oct 2014	Julie Wu
BZ#1132792 - Added a note on registering your system from the subscription manager GUI.		
BZ#1066161 - Added a note on network out of sync conditions.		
Revision 3.5-25	Fri 17 Oct 2014	Tahlia Richardson
BZ#1148210 - Updated version numbers in Self-Hosted engine topics.		
BZ#1149922 - Added RHEL 7 instructions to the procedure for checking network connectivity on a newly installed host.		
Revision 3.5-24	Wed 15 Oct 2014	Julie Wu
Updated the hardware certification link to https://hardware.redhat.com/ .		
BZ#1152523 - Added an important note on Hypervisors configured with bond or bridge device must be manually added from the Manager's side.		
Revision 3.5-23	Mon 13 Oct 2014	David Ryan
BZ#1066464 - Updated minimum system requirements.		
BZ#1151880 - Corrected spelling errors.		
Revision 3.5-22	Thu 09 Oct 2014	David Ryan
BZ#1150951 - Corrected product name errors.		
Revision 3.5-21	Thu 09 Oct 2014	David Ryan
BZ#1147711 - Correcting shutdown command syntax.		
Revision 3.5-20	Wed 08 Oct 2014	Julie Wu
BZ#1124129 - Included support for JSON protocol.		
Revision 3.5-19	Wed 08 Oct 2014	Lucy Bopf
BZ#1122596 - Updated the chapter on Red Hat Enterprise Linux hosts to include installation and configuration topics that previously appeared in a separate chapter.		
Revision 3.5-18	Wed 01 Oct 2014	Julie Wu
BZ#1145040 - Added an important note referring to the RHEL Security Guides.		
Removed all Beta references		
BZ#1147294 - Checked throughout the guide for outdated RHEL host versions.		
Revision 3.5-17	Fri 19 Sep 2014	Tahlia Richardson
BZ#1143843 - Removed RHN Classic references and replaced "Red Hat Network" with "Content Delivery Network."		
BZ#1094766 - Added a note about squid proxy connection timeout.		
BZ#1121013 - Added an Important box reminding the user to enable virtualization in the host's BIOS settings.		
Revision 3.5-16	Thu 18 Sep 2014	Andrew Burden
Brewing for 3.5-Beta.		
Revision 3.5-15	Wed 17 Sep 2014	Julie Wu
BZ#1142549 - Updated the affected sections with 3.5 beta channels and support for RHEL7 hosts.		
Revision 3.5-14	Thu 11 Sep 2014	Laura Novich

[BZ1063951](#) - Removed Step 4 from 7.5.3.5. Configuring Network Interfaces.

Revision 3.5-13	Thu 11 Sep 2014	Laura Novich
BZ1132792 - Removed instructions for installation via RHN Classic.		
Revision 3.5-12	Tue 09 Sep 2014	Julie Wu
Building for splash page.		
Revision 3.5-11	Mon 08 Sep 2014	Lucy Bopf
BZ#1123246 - Updated section on approving a Hypervisor.		
Revision 3.5-10	Mon 01 Sep 2014	Lucy Bopf
BZ#1123246 - Added a procedure for manually adding a Hypervisor host to the Manager.		
Revision 3.5-9	Thu 28 Aug 2014	Andrew Dahms
BZ#1083382 - Added a section outlining how to modify the Red Hat Enterprise Virtualization Hypervisor ISO file.		
BZ#853119 - Added a description of how to modify user and group IDs in the Red Hat Enterprise Virtualization Hypervisor ISO file.		
Revision 3.5-8	Mon 25 Aug 2014	Julie Wu
BZ#1123739 - Updated kbase article link for offline installation.		
Revision 3.5-7	Thu 21 Aug 2014	Tahlia Richardson
BZ#1122345 - Updated Red Hat Enterprise Virtualization installation instructions for 3.5-beta.		
BZ#1105691 - Added the required channels to the procedure.		
BZ#1123200 - Changed "select" to "enter" in the Other Device Selection section of the procedure or installing the Hypervisor interactively.		
Revision 3.5-6	Thu 21 Aug 2014	Lucy Bopf
BZ#1123226 - Revised the procedure for configuring networking in the Hypervisor to include IPv6 options.		
Revision 3.5-5	Wed 20 Aug 2014	Lucy Bopf
BZ#1122349 - Added new topics for 3.5-beta upgrade procedure.		
BZ#1123199 - Removed reference to the Hypervisor rebooting after registration with the Red Hat Enterprise Virtualization Manager.		
BZ#1123212 - Revised the Hypervisor install procedure to reflect the actual output in the Data field during storage setup.		
BZ#1123214 - Added a note that the Hypervisor will accept a weak password.		
BZ#1123216 - Removed a reference to a message that no longer appears after a successful Hypervisor installation.		
BZ#1123235 - Revised the procedure for registering the Hypervisor to the Manager to exclude the 'Retrieve Certificate' button.		
BZ#1121854 - Changed the default answer for confirming update of hosts in a self-hosted environment from 'no' to 'yes'.		
Revision 3.5-4	Wed 30 Jul 2014	Andrew Dahms

[BZ#1074917](#) - Added a note to the section on configuring a SPICE proxy regarding compatibility with SPICE HTML5 and noVNC.

[BZ#1044852](#) - Revised the procedure on installing the Red Hat Enterprise Virtualization Hypervisor to improve clarity.

Revision 3.5-3	Wed 23 Jul 2014	Lucy Bopf
-----------------------	------------------------	------------------

[BZ#1093486](#) - Removed the procedure for checking kvm module, and added a note about enabling virtualization in the BIOS.

[BZ#1114787](#) - Updated links to access.redhat.com to exclude '/site'.

Revision 3.5-2	Tue 15 Jul 2014	Andrew Burden
-----------------------	------------------------	----------------------

[BZ#1104114](#) - 'Installing the Self-Hosted Engine' now clearly lists the channels required to install the ovirt-self-hosted package.

Revision 3.5-1	Thu 5 Jun 2014	Lucy Bopf
-----------------------	-----------------------	------------------

Initial creation for the Red Hat Enterprise Virtualization 3.5 release.