

Red Hat Enterprise Linux 6

Hypervisor Deployment Guide

The complete guide to obtaining, deploying, configuring, and maintaining the Red Hat Enterprise Virtualization Hypervisor.

External Beta 1



Red Hat Enterprise Linux 6 Hypervisor Deployment Guide

The complete guide to obtaining, deploying, configuring, and maintaining the Red Hat Enterprise Virtualization Hypervisor.:

External Beta 1

Edition 1

Copyright © 2011 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at <http://creativecommons.org/licenses/by-sa/3.0/>. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, MetaMatrix, 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.

All other trademarks are the property of their respective owners.

1801 Varsity Drive
Raleigh, NC 27606-2072 USA
Phone: +1 919 754 3700
Phone: 888 733 4281
Fax: +1 919 754 3701

The Red Hat Enterprise Virtualization Hypervisor is a fully featured virtualization platform for quick, easy deployment and management of virtualized guests. The Hypervisor is designed for management via the Red Hat Enterprise Virtualization Manager. This *Hypervisor Deployment Guide* documents the steps required to obtain, deploy, configure, and maintain the Red Hat Enterprise Virtualization Hypervisor.

Preface	v
1. About this Guide	v
1.1. Audience	v
1.2. Documentation Suite	v
2. Document Conventions	vi
2.1. Typographic Conventions	vi
2.2. Pull-quote Conventions	viii
2.3. Notes and Warnings	viii
3. We Need Feedback!	ix
1. Introduction	1
2. Requirements	3
2.1. Hypervisor Requirements	3
2.2. Guest Requirements and Support Limits	6
2.3. Guest Operating System Support	7
3. Preparing Red Hat Enterprise Virtualization Hypervisor Installation Media	9
3.1. Preparation Instructions	9
3.2. Deploying Hypervisors with PXE and tftp	10
3.2.1. Booting a Hypervisor with PXE	11
3.3. Preparing a Hypervisor USB Storage Device	12
3.3.1. Making a USB Storage Device into a Hypervisor Boot Device	12
3.3.2. Booting a Hypervisor USB Storage Device	16
3.4. Preparing a Hypervisor from a CD-ROM or DVD	17
3.4.1. Making a Hypervisor CD-ROM Boot Disk	17
3.4.2. Booting a Hypervisor CD-ROM	17
4. Installation	19
4.1. Interactive Installation	19
4.1.1. Booting from the Installation Media	19
4.1.2. Installation Procedure	23
4.2. Automated Installation	26
4.2.1. How the Kernel Arguments Work	26
4.2.2. Required Parameters	26
4.2.3. Storage Parameters	26
4.2.4. Networking Parameters	29
4.2.5. Red Hat Network (RHN) Parameters	31
4.2.6. Authentication Parameters	33
4.2.7. Other Parameters	34
4.2.8. Example: Automated Hypervisor Installation	35
5. Configuration	37
5.1. Logging In	37
5.2. Status	37
5.3. Network	37
5.4. Security	40
5.5. Logging	40
5.6. Kernel Dump	41
5.7. Remote Storage	42
5.8. RHEV-M	42
5.9. Red Hat Network	43
6. Upgrading Red Hat Enterprise Virtualization Hypervisors	45
6.1. Upgrading a Hypervisor with the Manager	45
6.2. Upgrading a Red Hat Enterprise Virtualization Hypervisor with local media	46
6.3. Re-installing Hypervisors with the Manager	47

A. Security topics	49
B. Filesystem layout	51
C. Uninstallation	53
D. Revision History	55

Preface

This is a guide to the installation and configuration of Red Hat Enterprise Virtualization Hypervisors. The guide also provides step-by-step procedures to connect the Hypervisor with the Red Hat Enterprise Virtualization Manager. Advanced options are covered to assist users with configuration of Hypervisors in a wide variety of environments.

1. About this Guide

This guide describes the procedures for installation and configuration of Red Hat Enterprise Virtualization Hypervisors. Having read this guide you will be able to:

- create Hypervisor boot media,
- perform interactive installation of the Hypervisor,
- perform automated, or unattended, installation of the Hypervisor,
- configure the Hypervisor,
- attach the Hypervisor to a Red Hat Enterprise Virtualization Manager installation, and
- upgrade the Hypervisor as new versions become available.

Installation and configuration of the Red Hat Enterprise Virtualization Manager, other than the attachment of Hypervisors to the manager, is outside the scope of this document. For instruction on installation and configuration of the Red Hat Enterprise Virtualization Manager consult the *Red Hat Enterprise Virtualization Installation Guide*.

1.1. Audience

This guide is intended for use by those who need to install, configure, and maintain instances of the Red Hat Enterprise Virtualization Hypervisor. A relative level of comfort in the administration of computers that run Linux based operating systems would be beneficial but is not strictly required.

1.2. Documentation Suite

The Red Hat Enterprise Virtualization documentation suite provides information on installation, development of applications, configuration and usage of the Red Hat Enterprise Virtualization platform and its related products.

- *Red Hat Enterprise Virtualization Manager Release Notes* contain release specific information for Red Hat Enterprise Virtualization Managers.
- *Red Hat Enterprise Virtualization Quick Start Guide* provides quick and simple instructions for first time users to set up a basic Red Hat Enterprise Virtualization environment.
- *Red Hat Enterprise Virtualization Installation Guide* describes the installation prerequisites and procedures. Read this if you need to install Red Hat Enterprise Virtualization. The installation of hosts, manager and storage are covered in this guide. You will need to refer to the *Red Hat Enterprise Virtualization Administration Guide* to configure the system before you can start using the platform.
- *Red Hat Enterprise Virtualization Administration Guide* describes how to setup, configure and manage Red Hat Enterprise Virtualization. It assumes that you have successfully installed the Red Hat Enterprise Virtualization manager and hosts.

- *Red Hat Enterprise Virtualization Evaluation Guide* enables prospective customers to evaluate the features of Red Hat Enterprise Virtualization. Use this guide if you have an evaluation license.
- *Red Hat Enterprise Virtualization User Portal Guide* describes how users of the Red Hat Enterprise Virtualization system can access and use virtual desktops.
- *Red Hat Enterprise Virtualization Power User Portal Guide* describes how users of the Red Hat Enterprise Virtualization system can access and use virtual machines.
- *Red Hat Enterprise Linux Hypervisor Deployment Guide* (the book you are reading) describes how to deploy and install the hypervisor. Read this guide if you need advanced information about installing and deploying Hypervisors. The basic installation of Hypervisor hosts is also described in the *Red Hat Enterprise Virtualization Installation Guide*.
- *Red Hat Enterprise Linux V2V Guide* describes importing virtual machines from KVM, Xen and VMware ESX to Red Hat Enterprise Virtualization and KVM managed by libvirt.
- *Red Hat Enterprise Virtualization REST API Guide* describes how to use the REST API to set up and manage virtualization tasks. Use this guide if you wish to develop systems which integrate with Red Hat Enterprise Virtualization, using an open and platform independent API.
- *Red Hat Enterprise Virtualization Technical Reference Guide* describes the technical architecture of Red Hat Enterprise Virtualization and its interactions with existing infrastructure.

2. Document Conventions

This manual uses several conventions to highlight certain words and phrases and draw attention to specific pieces of information.

In PDF and paper editions, this manual uses typefaces drawn from the [Liberation Fonts](https://fedorahosted.org/liberation-fonts/)¹ set. The Liberation Fonts set is also used in HTML editions if the set is installed on your system. If not, alternative but equivalent typefaces are displayed. Note: Red Hat Enterprise Linux 5 and later includes the Liberation Fonts set by default.

2.1. Typographic Conventions

Four typographic conventions are used to call attention to specific words and phrases. These conventions, and the circumstances they apply to, are as follows.

Mono-spaced Bold

Used to highlight system input, including shell commands, file names and paths. Also used to highlight keycaps and key combinations. For example:

To see the contents of the file **my_next_bestselling_novel** in your current working directory, enter the **cat my_next_bestselling_novel** command at the shell prompt and press **Enter** to execute the command.

The above includes a file name, a shell command and a keycap, all presented in mono-spaced bold and all distinguishable thanks to context.

Key combinations can be distinguished from keycaps by the hyphen connecting each part of a key combination. For example:

¹ <https://fedorahosted.org/liberation-fonts/>

Press **Enter** to execute the command.

Press **Ctrl+Alt+F2** to switch to the first virtual terminal. Press **Ctrl+Alt+F1** to return to your X-Windows session.

The first paragraph highlights the particular keycap to press. The second highlights two key combinations (each a set of three keycaps with each set pressed simultaneously).

If source code is discussed, class names, methods, functions, variable names and returned values mentioned within a paragraph will be presented as above, in **mono-spaced bold**. For example:

File-related classes include **filesystem** for file systems, **file** for files, and **dir** for directories. Each class has its own associated set of permissions.

Proportional Bold

This denotes words or phrases encountered on a system, including application names; dialog box text; labeled buttons; check-box and radio button labels; menu titles and sub-menu titles. For example:

Choose **System** → **Preferences** → **Mouse** from the main menu bar to launch **Mouse Preferences**. In the **Buttons** tab, click the **Left-handed mouse** check box and click **Close** to switch the primary mouse button from the left to the right (making the mouse suitable for use in the left hand).

To insert a special character into a **gedit** file, choose **Applications** → **Accessories** → **Character Map** from the main menu bar. Next, choose **Search** → **Find...** from the **Character Map** menu bar, type the name of the character in the **Search** field and click **Next**. The character you sought will be highlighted in the **Character Table**. Double-click this highlighted character to place it in the **Text to copy** field and then click the **Copy** button. Now switch back to your document and choose **Edit** → **Paste** from the **gedit** menu bar.

The above text includes application names; system-wide menu names and items; application-specific menu names; and buttons and text found within a GUI interface, all presented in proportional bold and all distinguishable by context.

Mono-spaced Bold Italic or ***Proportional Bold Italic***

Whether mono-spaced bold or proportional bold, the addition of italics indicates replaceable or variable text. Italics denotes text you do not input literally or displayed text that changes depending on circumstance. For example:

To connect to a remote machine using ssh, type **ssh *username@domain.name*** at a shell prompt. If the remote machine is **example.com** and your username on that machine is john, type **ssh *john@example.com***.

The **mount -o remount *file-system*** command remounts the named file system. For example, to remount the **/home** file system, the command is **mount -o remount */home***.

To see the version of a currently installed package, use the **rpm -q *package*** command. It will return a result as follows: ***package-version-release***.

Note the words in bold italics above — username, domain.name, file-system, package, version and release. Each word is a placeholder, either for text you enter when issuing a command or for text displayed by the system.

Aside from standard usage for presenting the title of a work, *italics* denotes the first use of a new and important term. For example:

Publican is a *DocBook* publishing system.

2.2. Pull-quote Conventions

Terminal output and source code listings are set off visually from the surrounding text.

Output sent to a terminal is set in **mono-spaced roman** and presented thus:

```
books      Desktop  documentation  drafts  mss    photos  stuff  svn
books_tests Desktop1  downloads     images  notes  scripts svgs
```

Source-code listings are also set in **mono-spaced roman** but add syntax highlighting as follows:

```
package org.jboss.book.jca.ex1;

import javax.naming.InitialContext;

public class ExClient
{
    public static void main(String args[])
        throws Exception
    {
        InitialContext iniCtx = new InitialContext();
        Object          ref    = iniCtx.lookup("EchoBean");
        EchoHome        home   = (EchoHome) ref;
        Echo             echo   = home.create();

        System.out.println("Created Echo");

        System.out.println("Echo.echo('Hello') = " + echo.echo("Hello"));
    }
}
```

2.3. Notes and Warnings

Finally, we use three visual styles to draw attention to information that might otherwise be overlooked.



Note

Notes are tips, shortcuts or alternative approaches to the task at hand. Ignoring a note should have no negative consequences, but you might miss out on a trick that makes your life easier.



Important

Important boxes detail things that are easily missed: configuration changes that only apply to the current session, or services that need restarting before an update will apply. Ignoring a box labeled 'Important' will not cause data loss but may cause irritation and frustration.

**Warning**

Warnings should not be ignored. Ignoring warnings will most likely cause data loss.

3. We Need Feedback!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: <https://bugzilla.redhat.com/> against the component **Red Hat Enterprise Linux 6**.

When submitting a bug report, be sure to provide the following information:

- Manual's identifier: **doc-RHEV_Hypervisor_Deployment_Guide**
- Version number: **6**

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, include the section number and some of the surrounding text so we can find it easily.

Introduction

The Hypervisor is distributed as a compact image for use on a variety of installation media. It provides a minimal installation of Red Hat Enterprise Linux and includes the packages necessary to communicate with the Red Hat Enterprise Virtualization Manager.

The Hypervisor is certified for use with all hardware which has passed Red Hat Enterprise Linux certification except where noted in [Chapter 2, Requirements](#). The Hypervisor uses the Red Hat Enterprise Linux kernel and benefits from the default kernel's extensive testing, device support and flexibility.

Requirements

This chapter contains all system requirements and limitations which apply to Red Hat Enterprise Virtualization Hypervisors. These requirements are determined based on present hardware and software limits as well as testing and support considerations. System requirements and limitations will vary over time due to ongoing software development and hardware improvements.

2.1. Hypervisor Requirements

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

Table 2.1. Red Hat Enterprise Virtualization Hypervisor Requirements and Supported Limits

Item	Support Limit
CPU	<ul style="list-style-type: none"> • A minimum of 1 physical CPU is required. All CPUs must support: <ul style="list-style-type: none"> • the Intel® 64 or AMD64 CPU extensions, and • the AMD-V™ or Intel VT® hardware virtualization extensions. • A maximum of 128 physical CPUs is supported.
RAM	<ul style="list-style-type: none"> • A minimum of 512 MB of RAM is required. • A minimum of an additional 512 MB for each virtual machine is recommended. The amount of RAM required for each guest varies depending on: <ul style="list-style-type: none"> • the guest operating system's requirements, • the guests' application requirements, and • memory activity and usage of guests. <p>Additionally KVM is able to over-commit physical RAM for virtualized guests. It does this by only allocating RAM for guests as required and shifting underutilized guests into swap.</p> <ul style="list-style-type: none"> • A maximum of 1 TB of RAM is supported.
Storage	<p>The minimum supported internal storage for a Hypervisor is the total of the following list:</p> <ul style="list-style-type: none"> • The root partitions require at least 512 MB of storage. • The configuration partition requires at least 8 MB of storage.

Item	Support Limit
	<ul style="list-style-type: none"> • The recommended minimum size of the logging partition is 2048 MB. • The data partition requires at least 256 MB of storage. Use of a smaller data partition may prevent future upgrades of the Hypervisor from the Red Hat Enterprise Virtualization Manager. By default all disk space remaining after allocation of swap space will be allocated to the data partition. • The swap partition requires at least 8 MB of storage. The recommended size of the swap partition varies depending on both the system the Hypervisor is being installed upon and the anticipated level of overcommit for the environment. Overcommit allows the Red Hat Enterprise Virtualization environment to present more RAM to guests than is actually physically present. The default overcommit ratio is 0.5. <p>The recommended size of the swap partition can be determined by:</p> <ul style="list-style-type: none"> • 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. <div> <p>Example 2.1. Calculating Swap Partition Size</p> <p>For a system with 8 GB of RAM this means the formula for determining the amount of swap space to allocate is:</p> $(8 \text{ GB} \times 0.5) + 4 \text{ GB} = 8 \text{ GB}$ </div> <p>Please note that these are the <i>minimum</i> storage requirements for Hypervisor installation. It is recommended to use the default allocations which use more storage space.</p>

Item	Support Limit
PCI Devices	<ul style="list-style-type: none"> At least one network controller is required with a recommended minimum bandwidth of 1 Gbps.



Important — Virtualization Extensions

When the Red Hat Enterprise Virtualization Hypervisor boots a message may appear:

```
Virtualization hardware is unavailable.
(No virtualization hardware was detected on this system)
```

This warning indicates the virtualization extensions are either disabled or not present on your processor. Ensure that the CPU supports the listed extensions and they are enabled in the system BIOS.

To check that processor has virtualization extensions, and that they are enabled:

- At the Hypervisor boot screen press any key and select the **Boot** or **Boot with serial console** entry from the list. Press **Tab** to edit the kernel parameters for the selected option. After the last kernel parameter listed ensure there is a **Space** and append the **rescue** parameter.
- Press **Enter** to boot into rescue mode.
- At the prompt which appears, determine that your processor has the virtualization extensions and that they are enabled by running this command:

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

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

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

```
# lsmod | grep kvm
```

If the output includes **kvm_intel** or **kvm_amd** then the **kvm** hardware virtualization modules are loaded and your system meets requirements.



Important — Fakeraid Devices are not Supported

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.

2.2. Guest Requirements and Support Limits

The following requirements and support limits apply to guests that are run on the Hypervisor:

Table 2.2. Virtualized Hardware

Item	Limitations
CPU	<ul style="list-style-type: none"> A maximum of 64 virtualized CPUs per guest is supported.
RAM	<p>Different guests have different RAM requirements. The amount of RAM required for each guest varies based on the requirements of the guest operating system and the load under which the guest is operating. A number of support limits also apply.</p> <ul style="list-style-type: none"> A minimum of 512 MB of virtualized RAM per guest is supported. Creation of guests with less than 512 MB of virtualized RAM while possible is not supported. A maximum of 256 GB of virtualized RAM per 64 bit guest is supported. A maximum of 4 GB of virtualized RAM per 32 bit guest is supported. Note that not all 32 bit operating systems are able to register an entire 4 GB of RAM.
PCI devices	<ul style="list-style-type: none"> A maximum of 32 virtualized PCI devices per guest is supported. A number of system devices count against this limit, some of which are mandatory. Mandatory devices which count against the PCI devices limit include the PCI host bridge, ISA bridge, USB bridge, board bridge, graphics card, and the IDE or VirtIO block device.

Item	Limitations
Storage	<ul style="list-style-type: none"> A maximum of 8 virtualized storage devices per guest is supported.

2.3. Guest Operating System Support

Supported Guests

Red Hat Enterprise Virtualization presently supports the following virtualized guest operating systems:

- Red Hat Enterprise Linux 3 (32 bit and 64 bit)
- Red Hat Enterprise Linux 4 (32 bit and 64 bit)
- Red Hat Enterprise Linux 5 (32 bit and 64 bit)
- Red Hat Enterprise Linux 6 (32 bit and 64 bit)
- Windows XP Service Pack 3 and newer (32 bit only)
- Windows 7 (32 bit and 64 bit)
- Windows Server 2003 Service Pack 2 and newer (32 bit and 64 bit)
- Windows Server 2008 (32 bit and 64 bit)
- Windows Server 2008 R2 (64 bit only)



Note — Server Virtualization Validation Program (SVVP)

The Red Hat Enterprise Virtualization Hypervisor has been SVVP validated on both AMD and Intel systems.

Para-virtualized driver support

The para-virtualized drivers (the VirtIO drivers) support the following operating systems and versions. The para-virtualized drivers increase the performance for a guest's block and network devices. The drivers are only supported in environments with a Red Hat Enterprise Virtualization Manager.

Table 2.3. Para-virtualized driver support

Guest operating system	Para-virtualized drivers
Red Hat Enterprise Linux 4.8 and newer (32 bit and 64 bit)	Block and network drivers
Red Hat Enterprise Linux 5.4 and newer (32 bit and 64 bit)	Block and network drivers
Red Hat Enterprise Linux 6.0 and newer (32 bit and 64 bit)	Block and network drivers
Windows XP	Block and network drivers
Windows 7 (32 bit and 64 bit)	Block and network drivers
Windows Server 2003 R2 (32 bit and 64 bit)	Block and network drivers
Windows Server 2008 (32 bit and 64 bit)	Block and network drivers



Note — Windows Hardware Quality Labs (WHQL)

The para-virtualized drivers for Windows operating systems have been WHQL certified.

Red Hat Enterprise Virtualization agent support

The Red Hat Enterprise Virtualization agent is available for the following operating systems and versions. The agent is only supported in environments with a Red Hat Enterprise Virtualization Manager. Certain management and reporting functionality is only available for guests which have the agent installed.

- Red Hat Enterprise Linux 5.6 and newer (32 bit and 64 bit)
- Red Hat Enterprise Linux 6.0 and newer (32 bit and 64 bit)
- Windows XP Service Pack 3 and newer (32 bit only)
- Windows 7 (32 bit and 64 bit)
- Windows Server 2003 Service Pack 2 and newer (32 bit and 64 bit)
- Windows Server 2008 (32 bit and 64 bit)
- Windows Server 2008 R2 (64 bit only)

Preparing Red Hat Enterprise Virtualization Hypervisor Installation Media

This chapter covers creating installation media and preparing your systems before installing a Red Hat Enterprise Virtualization Hypervisor.

This chapter covers installing Red Hat Enterprise Virtualization Hypervisors on a local storage device. This storage device is a removable USB storage device, an internal hard disk drive or solid state drive. Once the Hypervisor is installed, the system will boot the Hypervisor and all configuration data is preserved on the system.

3.1. Preparation Instructions

The *rhev-hypervisor* package is needed for installation of Hypervisors. The *rhev-hypervisor* package contains the Hypervisor CD-ROM image. The following procedure installs the *rhev-hypervisor* package.

Entitlements to the Red Hat Enterprise Virtualization Hypervisor (v.6 x86-64) channel must be available on your Red Hat Network account to download the Hypervisor image. The channel's label is **rhev-x86_64-h-6**.

Downloading and Installing the RPM Package

The Red Hat Enterprise Virtualization Hypervisor package contains additional tools for USB and PXE installations as well as the Hypervisor ISO image.

1. Download the latest version of the *rhev-hypervisor** package from Red Hat Network. The list of Hypervisor packages is located at <https://rhn.redhat.com/rhn/channels/PackageList.do?cid=12058>.
2. Install the RPM on a Red Hat Enterprise Linux system. You must log in as the root user and navigate to the location of the downloaded file to perform this step.

```
# yum localinstall rhev-hypervisor*.rpm
```

The Hypervisor ISO image is installed into the `/usr/share/rhev-hypervisor/` directory and named **rhev-hypervisor.iso**.

The **rhevh-iso-to-disk** and **rhevh-iso-to-pxeboot** scripts are installed to the `/usr/bin` directory.

BIOS Settings and Boot Process Troubleshooting

Before installing Red Hat Enterprise Virtualization Hypervisors it is necessary to verify the BIOS is correctly configured for the chosen installation method. Many motherboard and PC manufacturers disable different booting methods in the BIOS. Most BIOS chips boot from the following devices in order:

1. 3.5 inch diskette
2. CD-ROM or DVD device

3. Local hard disk

Many BIOS chips have disabled one or more of the following boot methods: USB storage devices, CD-ROMs, DVDs or network boot. To boot from your chosen method, enable the method or device and set that device as the first boot device in BIOS.

Most but not all motherboards support the boot methods described in this chapter. Consult the documentation for your motherboard or system to determine whether it is possible to use a particular boot method.



Warning — BIOS Settings Vary Between Manufacturers

BIOS settings vary between manufacturers. Any specific examples of BIOS settings may be inaccurate for some systems. Due to this inconsistency, it is necessary to review the motherboard or system manufacturer's documentation.

Confirm Hardware Virtualization Support

Verify that your system is capable of running the Red Hat Enterprise Virtualization Hypervisor. Hypervisors require that virtualization extensions are present and enabled in the BIOS before installation proceeds.

1. Boot the Hypervisor from removable media. For example, a USB stick or CD-ROM.
2. Once the Hypervisor boot prompt is displayed, enter the command:

```
: linux rescue
```

3. Once the Hypervisor boots, verify your CPU contains the virtualization extensions with the following command:

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

Output displays if the processor has the hardware virtualization extensions.

4. Verify that the KVM modules load by default:

```
# lsmod | grep kvm
```

If the output includes **kvm_intel** or **kvm_amd** then the kvm hardware virtualization modules are loaded and the system meets the requirements.

3.2. Deploying Hypervisors with PXE and tftp

This section covers installing Hypervisors over a network with PXE and tftp. Configuring the DHCP and tftp servers for PXE booting is outside the scope of this book.

1. Install the **rhev-hypervisor** package. Refer to [Downloading and Installing the RPM Package](#)
2. Create **mlinuz** and **initrd** images with **rhev-iso-to-pxeboot**:

```
# rhev-iso-to-pxeboot /usr/share/rhev-hypervisor/rhev-hypervisor.iso
```

Your pxeboot image is complete.

Copy tftpboot/ subdirectory to /tftpboot or a subdirectory of /tftpboot.
Set up your DHCP, TFTP and PXE server to serve /tftpboot/.../pxeboot.0

Note: The initrd image contains the whole CD ISO and is consequently very large. You will notice when pxebooting that initrd can take a long time to download. This is normal behaviour.

- The output of **rhev-iso-to-pxeboot** command is a directory called **tftpboot** that has the following files in it:

- **pxelinux.0**
- **pxelinux.cfg/default**
- **vmlinuz0**
- **initrd0.img**

- It is possible to import the **vmlinuz** and **initrd** files into PXE and tftp servers.

Import the files to the appropriate directory.

- The **pxelinux.cfg/default** file provides a template for configuring the PXE server to export the Hypervisor image:

```
DEFAULT pxeboot
TIMEOUT 20
PROMPT 0
LABEL pxeboot
    KERNEL vmlinuz0
    APPEND rootflags=loop initrd=initrd0.img root=live:/rhev-hypervisor.iso
        rootfstype=auto ro liveimg nomodeset check rootflags=ro
crashkernel=512M-2G:64M,2G-:128M
    elevator=deadline processor.max_cstate=1 install rhgb rd_NO_LUKS rd_NO_MD rd_NO_DM
ONERROR LOCALBOOT 0
```

PXE booted Hypervisors rely on the PXE server passing the MAC address of the PXE interface to the kernel. This is provided by using the `IPAPPEND 2` parameter.

Modify the templates as required for your environment.



Important — Value of *root* Must Match ISO Name

The `root=live:/rhev-hypervisor.iso` parameter in **pxelinux.cfg/default** is a default value. If the ISO file you are using has a name other than **rhev-hypervisor.iso** it must be passed when calling **rhev-iso-to-pxeboot**. For example, for the ISO file **rhev_hypervisor_6_2.iso** use the command **rhev-iso-to-pxeboot rhev_hypervisor_6_2.iso**. This will produce the correct parameter `root=live:/rhev_hypervisor_6_2.iso` in **pxelinux.cfg/default**.

3.2.1. Booting a Hypervisor with PXE

For network booting the network interface card must support PXE booting.

To boot a Hypervisor from a PXE server:

1. Enter your system's BIOS. On most systems, the key or combination of keys is prompted shortly after the system has power. Usually, this key is **delete**, **F1** or **F2**.
2. Enable network booting if network booting is disabled.
3. Set the network interface card as the first boot device.
4. Boot the system. If the PXE parameters are configured correctly an automated installation will begin.
5. Change or disable network booting after the Hypervisor is installed. This is to avoid overwriting the installation on each reboot (unless this is desired functionality) and to prevent certain security vulnerabilities.



Important — PXE Boot Interface

The network interface used for PXE boot installation must be same interface used to connect to the Manager.



Note — Kernel Parameters

For more information on the kernel parameters, refer to [Section 4.2, “Automated Installation”](#).

The Hypervisor is now be installed.

3.3. Preparing a Hypervisor USB Storage Device

The Hypervisor install from USB storage devices and solid state disks. However, the initial boot/install USB device must be a separate device from the installation target. Network booting with PXE and tftp provides the greatest flexibility and scalability. For environments where network restrictions prevent network booting, or for systems without PXE capable network interface cards, a local media installation such as CD-ROM or USB is necessary. Booting from USB storage devices is a useful alternative to booting from CD, for systems without CD-ROM drives.



Note — USB Boot Support

Not all systems support booting from a USB storage device. Ensure that your system's BIOS supports booting from USB storage devices before proceeding.

3.3.1. Making a USB Storage Device into a Hypervisor Boot Device

This section covers making USB storage devices which are able to be used to boot Hypervisors.

3.3.1.1. Using **rhev-iso-to-disk** to Create USB Install Media

The **rhev-iso-to-disk** command will install a Hypervisor onto a USB storage device. The **rhev-iso-to-disk** command is part of the *rhev-hypervisor* package. Devices created with this command are able to boot the Hypervisors on systems which support booting via USB.

The basic **rhev-iso-to-disk** command usage follows this structure:

```
# rhev-iso-to-disk image device
```

Where the *device* parameter is the partition name of the USB storage device to install to. The *image* parameter is a ISO image of the Hypervisor. The default Hypervisor image location is **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**. The **rhev-iso-to-disk** command requires devices to be formatted with the FAT or EXT3 file system.



Note — Partitions and **rhev-iso-to-disk**

rhev-iso-to-disk uses a FAT or EXT3 formatted partition or block device.

USB storage devices are sometimes formatted without a partition table, use `/dev/sdb` or similar device name.

When a USB storage device is formatted with a partition table, use `/dev/sdb1` or similar device name.

1. Install the **rhev-hypervisor** package. Refer to [Downloading and Installing the RPM Package](#)
2. Use the **rhev-iso-to-disk** command to copy the .iso file to the disk. The `--format` parameter formats the disk. The `--reset-mbr` initializes the Master Boot Record (MBR). The example uses a USB storage device named `/dev/sdc`.

Example 3.1. Use of **rhev-iso-to-disk**

```
# rhev-iso-to-disk --format --reset-mbr /usr/share/rhev-hypervisor/rhev-
hypervisor.iso /dev/sdc
Verifying image...
/usr/share/rhev-hypervisor/rhev-hypervisor.iso:  eccc12a0530b9f22e5ba62b848922309
Fragment sums: 8688f5473e9c176a73f7a37499358557e6c397c9ce2dafb5eca5498fb586
Fragment count: 20
Checking: 100.0%

The media check is complete, the result is: PASS.

It is OK to use this media.
Copying live image to USB stick
Updating boot config file
Installing boot loader
syslinux: only 512-byte sectors are supported
USB stick set up as live image!
```

The USB storage device (`/dev/sdc`) is ready to boot a Hypervisor.

3.3.1.2. Using dd to Create USB Install Media

The **dd** command can also be used to install a Hypervisor onto a USB storage device. Media created with the command can boot the Hypervisor on systems which support booting via USB. Red Hat Enterprise Linux provides **dd** as part of the *coreutils* package. Versions of **dd** are also available on a wide variety of Linux and Unix operating systems.

Windows users are able to obtain the **dd** command through installation of **Red Hat Cygwin**, a free Linux-like environment for Windows. Refer to [Procedure 3.2, “Using dd to Create USB Install Media on Systems Running Windows”](#) for instruction on the installation and use of **Red Hat Cygwin** to install the Hypervisor to a USB storage device.

The basic **dd** command usage follows this structure:

```
# dd if=image of=device
```

Where the *device* parameter is the device name of the USB storage device to install to. The *image* parameter is a ISO image of the Hypervisor. The default Hypervisor image location is **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**. The **dd** command does not make assumptions as to the format of the device as it performs a low-level copy of the raw data in the selected image.

Procedure 3.1. Using dd to Create USB Install Media

1. Install the **rhev-hypervisor** package. Refer to [Downloading and Installing the RPM Package](#)
2. Use the **dd** command to copy the .iso file to the disk. The example uses a USB storage device named **/dev/sdc**.

Example 3.2. Use of dd

```
# 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 — All Data on the Device Specified Will be Overwritten

The **dd** command will overwrite all data on the device specified for the *of* parameter. Any existing data on the device will be destroyed. Ensure that the correct device is specified and that it contains no valuable data before invocation of the **dd** command.

Result:

The USB storage device (**/dev/sdc**) is ready to boot a Hypervisor.

Procedure 3.2. Using dd to Create USB Install Media on Systems Running Windows

1. Access <http://www.redhat.com/services/custom/cygwin/> and click the **Red Hat Cygwin official installation utility** link. The **rhsetup.exe** executable will download.
2. As the Administrator user run the downloaded **rhsetup.exe** executable. The **Red Hat Cygwin** installer will display.

3. Follow the prompts to complete a standard installation of **Red Hat Cygwin**. The *Coreutils* package within the *Base* package group provides the **dd** utility. This is automatically selected for installation.
4. Copy the **rhev-hypervisor.iso** file downloaded from **Red Hat Network** to **C:\rhev-hypervisor.iso**.
5. As the Administrator user run **Red Hat Cygwin** from the desktop. A terminal window will appear.



Note — Run Red Hat Cygwin as Administrator

On the **Windows 7** and **Windows Server 2008** platforms it is necessary to 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 **cat /proc/partitions** to see the drives and partitions currently visible to the system.

Example 3.3. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor #blocks name
 8      0  15728640 sda
 8      1   102400 sda1
 8      2  15624192 sda2
```

7. Plug the USB storage device which is to be used as the media for the Hypervisor installation into the system. Re-run the **cat /proc/partitions** command and compare the output to that of the previous run. A new entry will appear which designates the USB storage device.

Example 3.4. View of Disk Partitions Attached to System

```
Administrator@test /
$ cat /proc/partitions
major minor #blocks name
 8      0  15728640 sda
 8      1   102400 sda1
 8      2  15624192 sda2
 8     16    524288 sdb
```

8. Use the **dd** command to copy the **rhev-hypervisor.iso** file to the disk. The 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 3.5. Use of **dd** Command Under **Red Hat Cygwin**

```
Administrator@test /
```

```
$ dd if=/cygdrive/c/rhev-hypervisor.iso of=/dev/sdb& pid=$!
```

The provided command starts the transfer in the background and saves the process identifier so that it can be used to monitor the progress of the transfer. Refer to the next step for the command used to check the progress of the transfer.



Warning — All Data on the Device Specified will be Overwritten

The **dd** command will overwrite all data on the device specified for the *of* parameter. Any existing data on the device will be destroyed. Ensure that the correct device is specified and that it contains no valuable data before invocation of the **dd** command.

9. Transfer of the ISO file to the USB storage device with the version of **dd** included with **Red Hat Cygwin** can take significantly longer than the equivalent on other platforms.

To check the progress of the transfer in the same terminal window that the process was started in send it the **USR1** signal. This can be achieved by issuing the **kill** in the terminal window as follows:

```
kill -USR1 $pid
```

10. When the transfer operation completes the final record counts will be displayed.

Example 3.6. Result of **dd** Initiated Copy

```
210944+0 records in
210944+0 records out
108003328 bytes (108 MB) copied, 2035.82 s, 53.1 kB/s

[1]+ Done dd if=/cygdrive/c/rhev-hypervisor.iso of=/dev/sdb
```

Result:

The USB storage device (/dev/sdb) is ready to boot a Hypervisor.

3.3.2. Booting a Hypervisor USB Storage Device

Booting a Hypervisor from a USB storage device is similar to booting other live USB operating systems. To boot from a USB storage device:

1. Enter the system's BIOS menu to enable USB storage device booting if not already enabled.
 - a. Enable USB booting if this feature is disabled.
 - b. Set booting USB storage devices to be first boot device.
 - c. Shut down the system.
2. Insert the USB storage device that contains the Hypervisor boot image.
3. Restart the system.

4. The Hypervisor will boot automatically.

If the Hypervisor is running, you must now initialize the local storage device. Refer to [Section 4.1.1, “Booting from the Installation Media”](#) for details.

3.4. Preparing a Hypervisor from a CD-ROM or DVD

It is possible to install the Hypervisor with a CD-ROM or DVD.

3.4.1. Making a Hypervisor CD-ROM Boot Disk

Burn the Hypervisor image to a CD-ROM with the **cdrecord** command. The **cdrecord** command is part of the *cdrecord* package which is installed on Red Hat Enterprise Linux by default.

1. Verify that the *cdrecord* package is installed on the system.

Example 3.7. Verify Installation of *cdrecord* Package

```
# rpm -q cdrecord
cdrecord-2.01-10.7.el5
```

If the package version is in the output the package is available.

If it is not listed, install *cdrecord*:

```
# yum install cdrecord
```

2. Insert a blank CD-ROM or DVD into your CD or DVD writer.
3. Record the ISO file to the disc. The *cdrecord* command uses the following:

```
cdrecord dev=device /iso/file/path/
```

This example uses the first CD-RW (`/dev/cdrw`) device available and the default Hypervisor image location, `/usr/share/rhev-hypervisor/rhev-hypervisor.iso`.

Example 3.8. Use of **cdrecord** Command

```
# cdrecord dev=/dev/cdrw /usr/share/rhev-hypervisor/rhev-hypervisor.iso
```

If no errors occurred, the Hypervisor is ready to boot. Errors sometimes occur during the recording process due to errors on the media itself. If this occurs insert another writable disk and repeat the command above.

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.

3.4.2. Booting a Hypervisor CD-ROM

For many systems, the default BIOS configuration boots from CD-ROM first. If booting from CD-ROM is disabled or is not the first boot device refer to [BIOS Settings and Boot Process Troubleshooting](#) and your manufacturers manuals for more information.

To boot from CD-ROM insert the Hypervisor CD-ROM and then restart the computer.

The Hypervisor will start to boot. If the Hypervisor does not start to boot your BIOS may not be configured to boot from CD-ROM first or booting from CD-ROM may be disabled.

If the Hypervisor is running, you must now initialize the local storage device. Refer to [Section 4.1.1, “Booting from the Installation Media”](#) for details.

Installation

This chapter documents the installation of the Red Hat Enterprise Virtualization Hypervisor. The *Red Hat Enterprise Virtualization Installation Guide* covers installation of a Red Hat Enterprise Virtualization Manager or a Red Hat Enterprise Linux host.

Red Hat Enterprise Virtualization Hypervisors are able to use Storage Area Networks (SANs) and other network storage for storing virtualized guest images. Hypervisors are also able to be installed on SANs, provided that the Host Bus Adapter (HBA) permits configuration as a boot device in BIOS.

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.



Important — Red Hat Enterprise Virtualization Manager Installation

The Red Hat Enterprise Virtualization Manager for the environment must be installed and configured before Red Hat Enterprise Virtualization Hypervisors. Refer to the *Red Hat Enterprise Virtualization Installation Guide* for instructions on installing the Manager.

There are two methods for installing Red Hat Enterprise Virtualization Hypervisors:

- Interactive Installation (see [Section 4.1, “Interactive Installation”](#)).
- Automated Installation with Kernel Parameters (see [Section 4.2, “Automated Installation”](#)).

4.1. Interactive Installation

Red Hat Enterprise Virtualization Hypervisors must be installed on physical servers, not virtual machines.


The instructions in this section cover installation on a single system. When deploying on multiple systems always remember to use unique hostnames and IP addresses to avoid networking conflicts.

4.1.1. Booting from the Installation Media

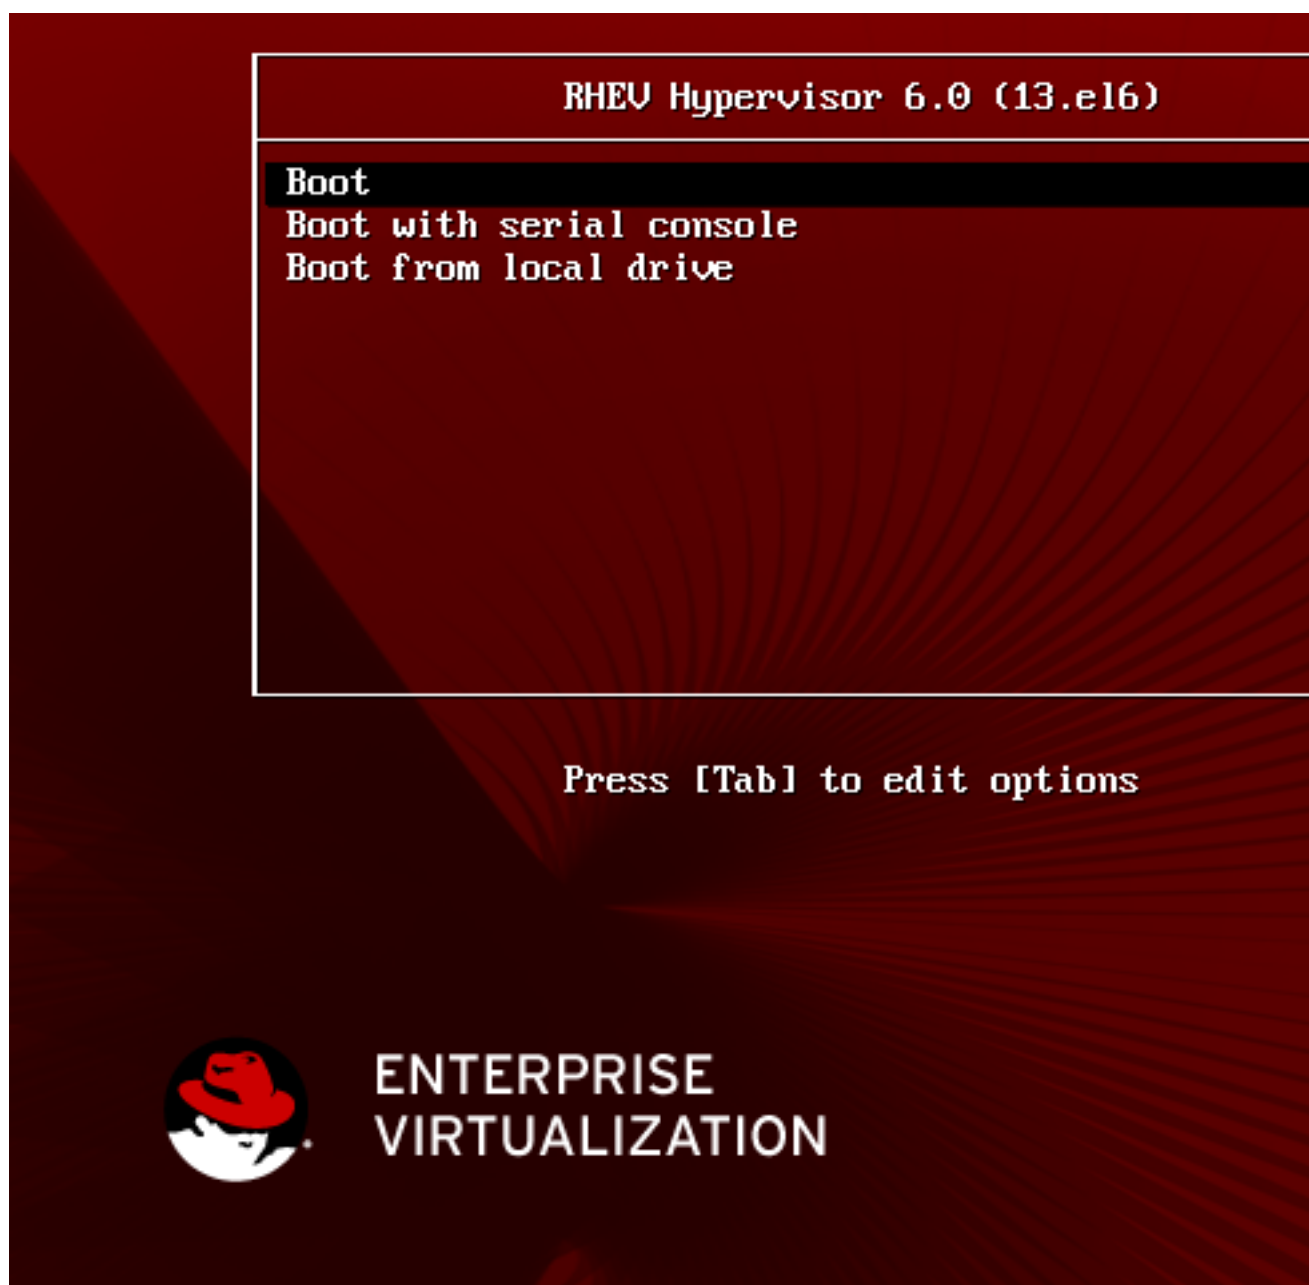
There are several methods for booting Hypervisors, refer to [Chapter 3, Preparing Red Hat Enterprise Virtualization Hypervisor Installation Media](#) for detailed instructions on preparing boot media for Red Hat Enterprise Virtualization Hypervisor installation.

Procedure 4.1. Booting from the Installation Media

1. Insert the Red Hat Enterprise Virtualization Hypervisor installation media.
2. Power on the system and ensure the system boots from the installation media.
3. The boot splash screen appears. If no input is provided, the Hypervisor installation will commence in 30 seconds, using default kernel parameters.



4. To modify the boot options, press any key. The boot menu will display.



The following boot options are available:

- **Boot**
Boot the Hypervisor installer.
- **Boot with Serial Console**
Boot the Hypervisor installer, with the console redirected to a serial device attached to `/dev/ttyS0`.
- **Boot from Local Drive**
Boot the operating system installed on the first local drive.

Select the appropriate boot option from the boot menu.

5. Where required additional kernel parameters should be appended to the default parameters displayed. A press of the **Enter** key boots the Hypervisor installation with the default kernel parameters. Alternatively press the **Tab** key to edit kernel parameters for the selected boot option.





Important — Kernel Parameters

In edit mode you are able to add or remove kernel parameters from the list. Kernel parameters must be separated from each other by a space. Once the desired kernel parameters have been set press **Enter** to boot the system. Alternatively pressing **Esc** reverts any changes that you have made to the kernel parameters.

For more information on the kernel parameters, refer to [Section 4.2, “Automated Installation”](#).



Note — Upgrading Existing Hypervisors

To upgrade an existing hypervisor installation, the kernel must be booted with the *upgrade* parameter. This will automatically upgrade and reboot the system, rather than displaying the interactive configuration menu. For more information, refer to [upgrade](#).



Note — PXE Booting

Kernel boot arguments are able to be appended to the PXE configuration file (`/pxelinux.cfg/default`). This file is able to be used to run an automated setup, covered in [Section 4.2, “Automated Installation”](#), and will be the more appropriate option in some environments.

4.1.2. Installation Procedure

When the Hypervisor is first booted the interactive installation script starts. This script facilitates installation of the Red Hat Enterprise Virtualization Hypervisor using graphical prompts. The following keys are used to manipulate the screens which support Hypervisor installation.

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.

Procedure 4.2. Hypervisor Installation

1. To commence Hypervisor installation select **Install Hypervisor** and press **Enter**.

2. **Disk Configuration**

The installation script automatically detects all disks attached to the system. This information is used to assist with selection of the boot and installation disks that the Hypervisor should use.

Each entry displayed on these screens indicates the **Location**, **Device Name**, and **Size (GB)** of the relevant disk.

- a. **Boot disk**

The first disk selection screen is used to select the disk from which the Hypervisor will boot. The Hypervisor's boot loader will be installed to the Master Boot Record (MBR) of the disk that is selected on this screen. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which you choose the boot device. Alternatively you are able to manually select a device, by specifying a block device name, by enabling the **Other Device** option.



Important — Boot Order

The disk selected must be identified as a boot device and appear in the boot order either in the system's BIOS or in a pre-existing boot loader.

Automatically Detected Device Selection

- i. Select the entry for the disk the Hypervisor is to boot from in the list.
- ii. Select the **<Continue>** button and press **Enter**. This action will save the boot device selection and start the next step of installation.

Manual Device Selection

- i. Select the **Other Device** entry from the list.
- ii. Select the **<Continue>** button and press **Enter**.
- iii. When prompted to **Please enter the disk to use for booting RHEV Hypervisor** enter the name of the block device from which the Hypervisor should boot.

Example 4.1. Other Device Selection

```
Please enter the disk to use for booting RHEV Hypervisor
/dev/sda
```

- iv. Select the **<Continue>** button and press **Enter**. This action will save the boot device selection and start the next step of installation.

Once a disk has been selected it is necessary to select the **<Continue>** button and press **Enter** to save the selection and continue with Hypervisor installation.

- b. **Installation Disk(s)**

The disk(s) selected for installation will be those to which the Hypervisor itself is installed. The Hypervisor attempts to automatically detect the disks attached to the system and presents the list from which installation devices are chosen.



Warning — Data Loss

All data on the selected storage device(s) will be destroyed.

- i. Select each disk which the Hypervisor is to use for installation and press **Space** to toggle it to enabled. Repeat this step for all disks you want the Hypervisor to use. Where other devices are to be used for installation, either solely or in addition to those which are listed automatically, enable the **Other Device** option.
- ii. Select the **<Continue>** button and press **Enter** to continue.
- iii. Where the **Other Device** option was specified a further prompt will appear. Enter the name of each additional block device to use for Hypervisor installation separated by a comma. Once all required disks have been selected then select the **<Continue>** button and press **Enter**.

Example 4.2. Other Device Selection

```
Please select the disk(s) to use for installation of RHEV Hypervisor
Enter multiple entries separated by commas
/dev/mmcblk0,/dev/mmcblk1_____
```

Once the installation disk, or disks, have been selected the next stage of the installation starts.

3. Password

The Hypervisor requires that a password be set to protect local console access by the `admin` user. The installation script prompts you to enter the desired password in both the **Password** and **Confirm Password** fields.

A strong password must be used. Strong passwords consist of a mix of uppercase, lowercase, numeric, and punctuation characters. They are six or more characters long and do not contain dictionary words.

Once a strong password has been entered select **<Install>** and press **Enter** to install the Hypervisor to disk.

Result:

Once installation is complete the message **RHEV Hypervisor Installation Finished Successfully** will be displayed. Select the **<Restart>** button and press **Enter** to reboot the system.

Further post installation configuration is required to connect the Hypervisor to the Red Hat Enterprise Virtualization Manager. See [Chapter 5, Configuration](#) for further details.



Note — Remove Boot Media

The boot media should be removed and the boot device order changed to prevent the installation sequence restarting after the system reboots.

4.2. 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 covered by [Section 4.1, “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.

4.2.1. 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** script 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 executes special Red Hat Enterprise Virtualization scripts and start up procedures.
3. An automated installation begins if all the required parameters are set.

4.2.2. Required Parameters

At a minimum, the following parameters are required for an automated installation:

1. **storage_init** to initialize a local storage device.
2. **BOOTIF** 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**.
3. **management_server** to specify the Manager server.

4.2.3. 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.

Presently, Hypervisors use one storage device for local installation. There are four 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* 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
```

- 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*, or *cciss* you also have the option of appending a serial number to explicitly set which device to use. The serial number for the device is determined by running the command shown in [Example 4.3, “Finding udev Serial Numbers”](#).

Example 4.3. Finding udev Serial Numbers

This command lists serial numbers for all disks attached to the system.

```
$ for d in /dev/sd?; do echo $d `udevadm info -q env -n $d | grep ID_SERIAL=`; done
/dev/sda ID_SERIAL=ST9500325AS_6VE867X1
```

When providing both a storage type and the serial number you should ensure that the two values are separated by a colon (:), for example:

```
storage_init=cciss:3600508b100104a3953545233304c0003
```



Note — Device Names are not Persistent

Consistency of devices 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*= there are six fields separated by colons. Not all fields have to be specified, those that you do not define during installation will be assigned their default value.

Chapter 4. Installation

The first and third values represent the boot and root partitions respectively, which have fixed sizes. These values cannot be set and should be left undefined.

All values are in megabytes (MB). Do not append units onto the end of the values.

Setting a size value of `-1` sets the partition to take up all remaining disk space. Note that this can only be used with the Data partition.

The following is the standard format of the `storage_vol` parameter with each element described in the list below.

Example 4.4. Format of the `storage_vol` Parameter

```
storage_vol=:SWAP::CONFIG:LOGGING:DATA
```

- *SWAP*

The swap partition is used for swapping pages of memory which 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.

Red Hat [Knowledgebase](#)¹ has an article on determining the size of the swap partition.

Use the formula from the Red Hat Knowledgebase 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 is required and the recommended size is 2048 MB.

- *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 the host's system RAM. The data partition also stores the Hypervisor ISO file for Hypervisor upgrades.

The data partition should be at least 1.5x as large as the RAM on the host system plus an additional 512 MB in size. The minimum size is 256 MB.

The default size for the data partition is the remaining available disk space (labeled as `-1`).

Example 4.5. Using the *storage_vol* Parameter to Partition Default Sizes

```
storage_vol=:::
```

Example 4.6. Using the *storage_vol* Parameter to Partition Certain sizes

The Boot partition is always omitted, and therefore defined as the fixed size of 50 MB.

The Swap partition is defined as 4000 MB.

The Root partition is always omitted, and therefore defined as the fixed size of 512 MB.

The Config partition is defined as 8 MB.

The Logging partition is defined as 2048 MB.

The Data partition is defined to take up all remaining disk space.

```
storage_vol=4000::8:2048:-1
```

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*, and
- optionally, a storage target name as specified by the naming authority - preceded by a colon.

Example 4.7. *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
```

4.2.4. 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.

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 — If Booting with PXE

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#** (where # is the number of the NIC) - indicates to use the NIC as determined by the kernel driver initialization order. 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.
- **<MAC>** - indicates to use the MAC address explicitly defined inside the brackets.

Example 4.8. To use the First NIC

```
BOOTIF=eth0
```

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 the word **dhcp** (for DHCP).

Example 4.9. Using an IP Address

```
ip=192.168.1.1
```

Example 4.10. Using DHCP

```
ip=dhcp
```

netmask

The *netmask* parameter sets the subnet mask for the IP address defined with the *ip* parameter.

Example 4.11. Using a Subnet Mask

```
netmask=255.255.255.0
```


gateway

The *gateway* parameter sets the Internet gateway.

Example 4.12. Setting the Gateway

```
gateway=192.168.1.246
```

dns

The *dns* parameter sets the addresses of one or more DNS servers. Each DNS server must be separated by a colon.

Example 4.13. Using two DNS Servers

```
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.

Example 4.14. Setting the Hypervisor *hostname*

```
hostname=rhev1.example.com
```

ntp

The *ntp* parameter sets the addresses of one or more Network Time Protocol servers. Each NTP server must be separated by a colon.

Example 4.15. Using two NTP Servers

```
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.

Example 4.16. Using a VLAN

```
vlan=VLAN-ID
```

4.2.5. 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 *rhn_activationkey* or both the *rhn_username* and *rhn_password* parameters must be provided. Where registration is to occur against a satellite server, the *rhn_url* parameter must be provided.

rh_n_username

The *rh_n_username* parameter sets the username used to connect to RHN.

Example 4.17. Setting a Username of testuser

```
rh_n_username=testuser
```

rh_n_password

The *rh_n_password* parameter sets the password used to connect to RHN.

Example 4.18. Setting a Password of testpassword

```
rh_n_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.

Example 4.19. Setting an Activation Key

```
rh_n_activationkey=7202f3b7d218cf59b764f9f6e9fa281b
```

rh_n_url

The *rh_n_url* parameter sets the URL of the satellite server used to register the host.

Example 4.20. Setting a Satellite Server URL

```
rh_n_url=https://satellite.example.com
```

rh_n_ca_cert

The *rh_n_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 *rh_n_url/pub/RHN-ORG-TRUSTED-SSL-CERT*

Example 4.21. Setting a CA Certificate URL

```
rh_n_ca_cert=https://satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT
```

rh_n_profile

The *rh_n_profile* parameter sets the name of the profile to be registered with RHN for this host. The default value is the system hostname.

Example 4.22. Setting a Profile Name of testhost

```
rh_n_profile=testhost
```

4.2.6. 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 — Setting the root Password is not Supported

The root password is not set by default and is not supported unless enabled at the request of Red Hat support.

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*.

Example 4.23. Enabling SSH Password Authentication

```
ssh_pwauth=1
```



Important — SSH Password Authentication is not Supported

SSH password authentication is disabled by default and is not supported unless enabled at the request of Red Hat support.

4.2.7. Other Parameters

firstboot

The *firstboot* parameter starts the interactive configuration menu. On systems that have a Red Hat Enterprise Virtualization Hypervisor installed and some systems with LVM volumes, the *firstboot* parameter may be required to trigger the interactive installation.

Combination of the *firstboot* parameter with the *storage_init* parameter erases data on the specified disk. The *reinstall* parameter is an alias for *firstboot*. The two are able to be used interchangeably.

install

The *install* parameter forces the installation menu to be displayed on boot even when the Hypervisor has previously been installed. The *install* parameter is intended to be used when booting from CD-ROM, DVD, USB, or PXE media.

local_boot

The *local_boot* parameter is an alias for the *upgrade* parameter.

netconsole_server

The *netconsole_server* parameter sets the address of a server to which kernel messages should be logged. The *netconsole_server* parameter takes an IP address or fully qualified domain name and, optionally, a port (the default port is 25285).

Example 4.24. Setting a NetConsole Server at `rhev.example.com:25285`

```
netconsole_server=rhev.example.com:25285
```

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 is required for an automated installation.

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). It is required for an automated installation.

Example 4.25. Connecting to a Red Hat Enterprise Virtualization Manager at `rhev.example.com: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.

Example 4.26. Changing the Memory Overcommit to 0.7

```
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 *gpxe* and *etherboot*. For compatibility with Red Hat Enterprise Virtualization Hypervisor 5.4-2.1, the default value is *etherboot*.

Example 4.27. Using the *gpxe* Boot Loader

```
qemu_pxe=gpxe
```

reinstall

The *reinstall* parameter starts the interactive configuration menu. On systems that have a Red Hat Enterprise Virtualization Hypervisor installed and some systems with LVM volumes, the *reinstall* parameter may be required to trigger the interactive installation.

Combination of the *firstboot* parameter with the *storage_init* parameter erases data on the specified disk. The *reinstall* parameter is an alias for *firstboot*. The two are able to be used interchangeably.

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. For further information on Hypervisor uninstallation see [Appendix C, Uninstallation](#).

4.2.8. Example: Automated Hypervisor Installation

This example uses the kernel command line parameters for an automated Hypervisor installation.

**Important — Customize Before Use**

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_server=rhevm.example.com
```



Note — For PXE Installations

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.

Configuration

5.1. Logging In

The Hypervisor allows local console logins to facilitate post-installation configuration. The login prompt used is displayed once the Hypervisor has booted:

```
Please login as 'admin' to configure the node
localhost login:
```

Type `admin` at the prompt and press **Enter**. When prompted enter the password which was set during the installation process and press **Enter** again to log in.

The Hypervisor configuration menu will then be displayed. The menu facilitates interactive configuration of the Hypervisor. Throughout the remainder of this chapter it will be referred to as the main menu. The main menu provides access to multiple screens which report on the status and configuration of the hypervisor. They also provide the ability to change the hypervisor configuration.

The configuration interface is similar to that of the installation script. The same keys are used to navigate the menu and associated screens. Refer to [Menu Actions](#) to review the list of possible actions.

5.2. Status

The status screen displays a brief overview of the current state of the Hypervisor. The information displayed consists of:

- the hostname,
- the current status of the network connection,
- the destination(s) of logs and reports, and
- the number of active virtual machines.

The status screen also provides a number of buttons to change the state of the Hypervisor. They are:

- **<Lock>**: Locks the Hypervisor. The username and password must be entered to unlock the Hypervisor.
- **<Restart>**: Restarts the Hypervisor.
- **<Power Off>**: Turns the Hypervisor off.

5.3. Network

The **Network** screen is used to configure:

- the Hypervisor's hostname,
- the DNS server(s) to use,
- the NTP server(s) to use, and

- the network interface to use.

Procedure 5.1. Hostname Configuration

1. To set or change the hostname select the **Hostname** field and enter the new hostname.
2. Select **<Apply>**, and press **Enter** to save changes to the hostname.

Result:

The hostname is updated.

Procedure 5.2. DNS Configuration

The Hypervisor supports the specification of one or more Domain Name System (DNS) servers to use when resolving host and domain names.

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 to use.
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 to use.
3. Select **<Apply>**, and press **Enter** to save changes to the DNS configuration.

Result:

The primary and secondary DNS servers queried by the Hypervisor are updated.

Procedure 5.3. NTP Configuration

The Hypervisor supports the specification of one or more Network Time Protocol (NTP) servers with which the Hypervisor should synchronize the system clock. It is important that the Hypervisor is synchronized with the same time source as the Red Hat Enterprise Virtualization Manager. This ensures accurate time keeping across the Red Hat Enterprise Virtualization environment.

1. To set or change the primary NTP server select the **NTP Server 1** field and enter the IP address or hostname of the new primary NTP server to use.
2. To set or change the secondary NTP server select the **NTP Server 2** field and enter the IP address or hostname of the new secondary NTP server to use.
3. Select **<Apply>**, and press **Enter** to save changes to the NTP configuration.

Result:

The primary and secondary NTP servers queried by the Hypervisor are updated.

Procedure 5.4. Network Interface Configuration

For each network interface detected the Hypervisor will display the:

- **Device**,
- **Status**,
- **Model**, and
- **MAC Address**.

At least one network interface must be configured before the Hypervisor is able to connect with the Red Hat Enterprise Virtualization Manager.

1. Device Identification

Select the network interface to be configured from the list and press **Enter**.

In some cases it may be unclear which physical device an entry in the list refers to. Where this is the case the Hypervisor is able to blink the physical device's network traffic lights to assist with identification. To make use of this facility select the entry from the list and, then select the **<Flash Lights to Identify>** button. Press **Enter** and, take note of which physical device's lights start blinking. The configuration screen for the selected device will be displayed.

2. IPv4 Settings

The Hypervisor supports both dynamic (DHCP), and static IPv4 network configuration.

Dynamic (DHCP) Network Configuration

Dynamic network configuration allows the Hypervisor to be dynamically assigned an IP address via **DHCP**. To enable dynamic IPv4 network configuration select the **DHCP** option under **IPv4 Settings** and press **Space** to toggle it to enabled.

Static Network Configuration

Static network configuration allows the Hypervisor to be manually assigned an IP address. To enable static IPv4 network configuration select the **Static** option under **IPv4 Settings** and press **Space** to toggle it to enabled.

Selection of the **Static** option enables the **IP Address**, **Netmask**, and **Gateway** fields. The **IP Address**, **Netmask**, and **Gateway** fields must be populated to complete static network configuration.

In particular it is necessary that:

- the **IP Address** is not already in use on the network,
- the **Netmask** matches that used by other machines on the network, and
- the **Gateway** matches that used by other machines on the network.

Where it is not clear what value should be used for the **IP Address**, **Netmask**, or **Gateway** field consult the network's administrator or consider a dynamic configuration.

Example 5.1. 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_____
```

3. IPv6 Settings

The Red Hat Enterprise Virtualization Manager does not currently support IPv6 networking. IPv6 networking must remain set to **Disabled**.

4. VLAN Configuration

If VLAN support is required then populate the **VLAN ID** field with the VLAN identifier for the selected device.

5. Save Network Configuration

Once all networking options for the selected device have been set the configuration must be saved.

- Select the **<Apply>** button and press **Enter** to save the network configuration.

- b. The **Confirm Network Settings** dialog box will appear. Ensure that the **Ok** button is selected and press **Enter** to confirm.

Result:

The **Network** screen is displayed. The device is listed as **Configured**.

5.4. Security

The **Security** screen is used to change the admin password for both local and remote access. SSH password authentication is also enabled or disabled via this screen.

Procedure 5.5. Change Security Configuration

1. Enable SSH Password Authentication

To enable SSH password authentication for remote access select the **Enable ssh password authentication** option and press **Space** to toggle it to enabled.



Important — SSH Password Authentication is not Supported

SSH password authentication is disabled by default and is not supported unless enabled at the request of Red Hat support.

2. Change admin Password

- a. Enter the desired admin password in the **Password** field. You should 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. Enter the desired admin password in the **Confirm Password** field. Ensure that the value entered in the **Confirm Password** field matches the value entered in the **Password** field exactly. Where this is not the case an error message will be displayed to indicate that the two values are different.

3. Select **<Apply>** and press **Enter** to save the security configuration.

Result:

The security configuration has been updated.

5.5. Logging

The Hypervisor creates and updates a number of log files. The **Logging** screen allows configuration of a daemon to automatically export these log files to a remote server.

Procedure 5.6. Change Logging Configuration

1. Logrotate Configuration

The **logrotate** utility simplifies the administration of log files. The Hypervisor uses **logrotate** to rotate logs when they reach a certain file size.

Log rotation involves renaming the current log(s) and starting new ones in their place. The **Logrotate Max Log Size** value set on the **Logging** screen is used to determine when a log should be rotated.

Enter the **Logrotate Max Log Size** in kilobytes. The default maximum log size is 1024 kilobytes.

2. **Rsyslog Configuration**

The **rsyslog** utility is a multithreaded syslog daemon. The Hypervisor is able to use **rsyslog** to transmit log files over the network to a remote syslog daemon. For information on setting up the remote syslog daemon consult the *Red Hat Enterprise Linux — Deployment Guide*.

- a. Enter the remote **Rsyslog** server address in the **Server Address** field.
- b. Enter the remote **Rsyslog** server port in the **Server Port** field. The default port is **514**.

3. **netconsole Configuration**

The **netconsole** module allows kernel messages to be sent to a remote machine. The Hypervisor uses **netconsole** to transmit kernel messages over the network.

- a. Enter the **Server Address**.
- b. Enter the **Server Port**. The default port is **6666**.

4. **Save Configuration**

To save the logging configuration select **<Apply>** and press **Enter**.

Result:

The logging configuration has been updated and logs will be exported to the remote **Rsyslog** server specified.

5.6. Kernel Dump

Red Hat Enterprise Virtualization Hypervisor hosts generate a kernel dump (a **kdump** file) in the event of a system failure. These **kdump** files are essential for debugging and support.

The Hypervisor supports the export of kernel dumps by **kdump** using NFS or SSH so that they may be analyzed at a later date. Alternatively the hypervisor is able to store the kernel dumps locally. The **Kernel Dump** screen provides for configuration of this facility.

Procedure 5.7. **kdump** Configuration

1. **Kernel Configuration**

Crash dumps generated by **kdump** are exported over NFS or SSH. Select the desired transfer method and press **Space** to enable it. Alternatively, to keep the kernel dumps on the hypervisor's local storage, select **Restore (local)** and press **Space** to enable it.

Where export of the kernel dumps over NFS or SSH is chosen a location to which the **kdump** files will be exported to must also be specified.

a. **NFS location**

Set the NFS location to which crash logs should be exported in the **NFS Location** field. The **NFS Location** should be the full NFS path which includes fully qualified domain name and directory path.

Example 5.2. NFS Location

```
example.redhat.com:/var/crash
```

b. **SSH location**

Set the SSH location to which crash logs should be exported in the **SSH Location** field. The **SSH Location** should be the full SSH login which includes the fully qualified domain name and username.

Example 5.3. SSH Location

```
root@example.redhat.com
```

2. **Save Configuration**

To save the configuration the user must select **<Apply>** and press **Enter**.

Result:

The **Kernel Dump** configuration has been updated and kernel dumps will be exported to the remote server(s) specified.

5.7. Remote Storage

The Hypervisor supports the use of a remote iSCSI initiator for storage. The iSCSI initiator to use is set from the **Remote Storage** screen.

Procedure 5.8. Remote Storage Configuration

1. **iSCSI Initiator Name**

Enter the initiator name in the **iSCSI Initiator Name** field. 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*, and
- optionally, a storage target name as specified by the naming authority - preceded by a colon.

Example 5.4. iSCSI Initiator Name

```
iqn.2011-08.com.redhat.demo:target1
```

2. **Save Configuration**

To save the configuration the user must select **<Apply>** and press **Enter**.

Result:

The **Remote Storage** configuration has been updated.

5.8. RHEV-M

To configure the Hypervisor to connect to the Red Hat Enterprise Virtualization Manager the details of the server on which it resides must be provided. The settings entered on the Hypervisor must match those that were used during installation of the Red Hat Enterprise Virtualization Manager.

Procedure 5.9. RHEV-M Configuration**1. Enter Management Server details**

- 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**. Where a different port was selected during Red Hat Enterprise Virtualization Manager installation then it should be specified here, replacing the default value.

2. Save Configuration

To save the configuration the user must select **<Apply>** and press **Enter**.

Result:

The **RHEV-M** configuration has been updated.

5.9. Red Hat Network

The Hypervisor is able to retrieve security updates from either the Red Hat Network or a Satellite server. Where necessary the Hypervisor is also able to connect to the remote updates server via a HTTP proxy.

Procedure 5.10. Register with Red Hat Network**1. Authentication**

Enter your Red Hat Network username in the **Login** field.

Enter your Red Hat Network password in the **Password** field.

2. Profile Name

Enter the profile name to be used for the system in the **Profile Name** field. This is the name that the system will appear under when viewed via the Red Hat Network.

3. Update Source

The Hypervisor is able to receive updates either directly from the Red Hat Network or, if available, via a Satellite installation.

To Connect Directly to RHN

Select the **RHN** option and press **Space** to toggle it to enabled. The **URL** and **CA** values do not need to be provided.

Example 5.5. Red Hat Network Configuration

```
[*] RHN [ ] Satellite
URL: _____
CA : _____
```

To Connect via Satellite

- a. Select the **Satellite** option and press **Space** to toggle it to enabled.
- 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** field.

Example 5.6. Satellite Configuration

```
[ ] RHN [*] Satellite
```

```
URL:  https://your-satellite.example.com
CA :  https://your-satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT
```

4. HTTP Proxy

Where a HTTP proxy is in use the details to connect to it must be provided. To connect to the Red Hat Network or a Satellite server via a proxy you must enter:

- the proxy **Server's** network address,
- the **Port** to connect to the proxy on, and
- optionally, the **Username** and **Password** to use to connect to the proxy.

In environments where a HTTP proxy is not in use it is safe to ignore this step.

Example 5.7. HTTP Proxy Configuration

```
HTTP Proxy
Server: proxy.example.com__ Port: 8080_
Username: puser_____ Password: *****_____
```

5. Save Configuration

To save the configuration the user must select **<Apply>** and press **Enter**.

Result:

The **Red Hat Network** configuration has been updated.

Upgrading Red Hat Enterprise Virtualization Hypervisors

Red Hat Enterprise Virtualization Hypervisors can be updated to get the latest features, bug fixes and security patches.

6.1. Upgrading a Hypervisor with the Manager

You can upgrade and reinstall a Red Hat Enterprise Virtualization Hypervisor host from an ISO image stored on the Red Hat Enterprise Virtualization Manager. Upgrading and reinstalling means that you are stopping and restarting the host. Virtual machines are automatically migrated to a different host, however it is recommended that the upgrade is performed at a time when usage of the system is at its lowest. Ensure that the cluster contains more than one host before performing an upgrade.

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



Warning — Potential Data Loss

Upgrading Hypervisor hosts involves shutting down, deactivating guests, and restarting the physical server. If any virtual machines are running on the Hypervisor, all data and configuration details may be destroyed if they are not shut down. Upgrading Hypervisors must be carefully planned and executed with care and consideration.

Prerequisites

Before upgrading a Hypervisor:

1. Download the latest Red Hat Enterprise Virtualization Hypervisor package from Red Hat Network. It is available at <https://rhn.redhat.com/rhn/channels/PackageList.do?cid=12058>. Install the package on the Red Hat Enterprise Virtualization Manager server.



Important — One Host Must Remain Active

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

Procedure 6.1. Upgrading a Red Hat Enterprise Virtualization Hypervisor

1. Click the **Hosts** tab. A list of hosts displays. Select the host that you intend to upgrade.

If the host is not displayed, or the list of hosts is too long to filter visually, perform a search to locate the host.

2. On the **Details** pane, click the **General** Tab.

An Alert Message indicates that a new version of the Red Hat Enterprise Virtualization Hypervisor is available. The **Upgrade** link is disabled if the host has a status of **Up**. A tooltip directs you to switch to maintenance mode to enable upgrade.

3. Select the host and click the **Maintenance** button. This will cause any VMs running on the host to be migrated to other hosts. If the host is the SPM, this function will be moved to another host. The status of the host changes as it enters maintenance mode. When the host status is **Maintenance**, the message in the general tab changes, providing you with a link which when clicked will re-install or upgrade the host.
4. Click the link. The **Install Host** dialog displays.
5. Select the appropriate file from the list of available ISOs to upgrade the host. This is usually **rhev-hypervisor.iso**.
6. Click **OK** to upgrade and re-install the host. The dialog closes, the details of the host are updated in the **Hosts** tab, and the status changes appropriately.

The host status transitions through the following stages: **Installing**, **Reboot**, **Non Responsive**, and **Up**. These are all expected, and each stage will take some time.

7. Once successfully upgraded, the host displays a status of **Up**. Any virtual machines that were migrated off the host, are at this point able to be migrated back onto the upgraded host.

6.2. Upgrading a Red Hat Enterprise Virtualization Hypervisor with local media

It is possible to upgrade to new versions of the Hypervisor using local media. Before commencing the upgrade process it is necessary to prepare the local media with the latest version of the Hypervisor.

- CD-ROM or DVD media must be prepared as described in [Section 3.4.1, “Making a Hypervisor CD-ROM Boot Disk”](#).
- USB boot media must be prepared as described in [Section 3.3.1, “Making a USB Storage Device into a Hypervisor Boot Device”](#).

Once prepared the local media is used to boot the Hypervisor with the **upgrade** parameter as described in [Procedure 6.1, “”](#).

1. Insert the installation media, and start the system.
2. Start the system. When the automatic boot prompt appears, press **Enter**.

Automatic boot in 30 seconds...

3. Select the **boot** option, and press **Tab** to enter edit mode.
4. The kernel parameters in use are displayed. Append the **upgrade** parameter to the list.

```
vmlinux0 initrd=initrd0.img root=live:CDLABEL=rhev-hypervisor rootfstype=auto ro
liveimg nomodeset check rootflags=ro crashkernel=512M-2G:64M,2G-:128M elevator=deadline
processor.max_cstate=1 install rhgb rd_NO_LUKS rd_NO_MD rd_NO_DM upgrade
```

Result:

The Red Hat Enterprise Virtualization Hypervisor is upgraded.

6.3. Re-installing Hypervisors with the Manager

Re-installing Red Hat Enterprise Virtualization Hypervisors is the same procedure as upgrading, refer to [Procedure 6.1, “Upgrading a Red Hat Enterprise Virtualization Hypervisor”](#) for details on upgrading.

To reset settings, remove the existing Hypervisor installation, as documented in [Appendix C, Uninstallation](#), and then re-install the Hypervisor.

Appendix A. Security topics

The Red Hat Enterprise Virtualization Hypervisor has various security features enabled. Security-Enhanced Linux (SELinux) and the **iptables** firewall are fully configured and on by default.

Administrators can receive the latest security advisories from the Red Hat Enterprise Virtualization watch list. Subscribe to the Red Hat Enterprise Virtualization watch list to receive new security advisories for Red Hat Enterprise Virtualization products by email. Subscribe by completing this form: <http://www.redhat.com/mailman/listinfo/rhev-watch-list/>.

Red Hat Enterprise Virtualization uses various network ports for management and other virtualization features. These ports must be open for Red Hat Enterprise Linux to function as a host with Red Hat Enterprise Virtualization. The list below covers ports and their usage by Red Hat Enterprise Virtualization:

- ICMP requests must be accepted. ICMP packets are used for network testing by the Manager.
- Port 22 should be open for SSH access and the initial installation.
- Ports 80 or 443 (depending on the security settings on the Manager) are used by the `vdsm-reg` service to communicate information about the host.
- Ports 5634 to 6166 are used for guest console access.
- Ports 49152 to 49216 are used for migrations. Migration may use any port in this range depending on the number of concurrent migrations occurring.
- Port 54321 is used by default, by *VDSM* for management, storage and inter-host communication. This port can be modified.

Appendix B. Filesystem layout

This Appendix provides an over view of the filesystem layout used by the Red Hat Enterprise Virtualization Hypervisor. The hypervisor directory layout provides a number of top level directories in addition to those provided by a standard Red Hat Enterprise Linux installation.

/config

The **/config** directory contains all persistent configuration files for the Red Hat Enterprise Virtualization Hypervisor. These files control passwords, storage configuration, security and networking.

The **/config** directory must be at least 8MB.

/boot

The **/boot** directory contains the boot loader, the kernel and the **initramfs** file.

/liveos

The **/liveos** directory contains a compressed Red Hat Enterprise Virtualization Hypervisor live CD image. The Red Hat Enterprise Virtualization Hypervisor boots and runs from the ISO image file in this directory.

The **/liveos** directory is not normally visible on the running system. This is the folder containing the CD-ROM ISO. During an upgrade **/dev/HostVG/Root** is temporarily mounted to **/liveos**

/var/log

Contains all the logs for the Hypervisor.

The **log** directory must be at least 2048MB. The default size of the **log** directory is 2048MB.

/var/log/core

Contains core dumps from the Hypervisor which can be used for debugging and support.

/var/vdsm/

The **/var/vdsm/** is used by the **vdsm** daemon for storing volatile data, including **/var/vdsm/ts** which stores **vdsm truststore**.

/rhev/data-center

Contains links to Storage Domains.

/data

This directory contains virtual machine cache data and other miscellaneous files.

The data partition must be at least as large as the RAM on the host system plus an additional 512MB in size. A data partition at least one and a half times as large as the RAM on the host system is recommended.

The default size for the data partition is the remaining available disk space.

Appendix C. Uninstallation

This appendix covers uninstallation of the Red Hat Enterprise Virtualization Hypervisors from local storage devices.



Warning — Data loss

All data on the selected storage device(s) will be destroyed.

Procedure C.1. Removing Red Hat Enterprise Virtualization Hypervisors

1. Boot the Hypervisor. Refer to [Chapter 3, Preparing Red Hat Enterprise Virtualization Hypervisor Installation Media](#) for details on booting Red Hat Enterprise Virtualization Hypervisors.
2. Start the system. When the automatic boot prompt appears, press **Enter**.

```
Automatic boot in 30 seconds...
```

3. Select the **boot** option, and press **Tab** to enter edit mode.
4. The kernel parameters in use are displayed. These are the parameters passed to the kernel when starting the Hypervisor.

```
vmlinux0 initrd=initrd0.img root=live:CDLABEL=rhev-hypervisor rootfstype=auto ro  
liveimg nomodeset check rootflags=ro crashkernel=512M-2G:64M,2G-:128M elevator=deadline  
processor.max_cstate=1 install rhgb rd_NO_LUKS rd_NO_MD rd_NO_DM
```

5. To perform uninstallation of the Hypervisor two parameters need to be added to the end of the **kernel** command string. The parameters which need to be added are:
 - the **uninstall** parameter, and
 - the **storage_init** parameter.

The **uninstall** parameter takes no arguments and specifies that the Hypervisor is to be uninstalled.

The **storage_init** parameters specifies the device the Hypervisor resides on. This information is also used to determine which device the Hypervisor should be removed from.

Once the two parameters have been added to the **kernel** string press **Enter**. This action saves the changes to the **kernel** string, for this boot only, and displays the previous screen.

Example C.1. Hypervisor Uninstallation

In this example the Hypervisor is installed to the device `/dev/sda`. The parameters that must be added to the **kernel** command string are:

```
uninstall storage_init=/dev/sda
```

6. Press **Enter** to boot the Hypervisor. The Hypervisor will uninstall itself immediately. Once the Hypervisor has been removed the system will reboot.

Result:

The Red Hat Enterprise Virtualization Hypervisor has been removed from the specified device.

Appendix D. Revision History

Revision 1-0 Thursday June 23 2010

Stephen Gordon sgordon@redhat.com

- Updated USB boot media creation instructions.
- Updated PXE boot media creation instructions.
- Updated installation instructions to match new UI.
- Updated configuration instructions to match new UI.
- Added documentation of new kernel parameters.

