Red Hat Enterprise Linux 5.7

Hypervisor Deployment Guide

The complete guide to installing, deploying and maintaining Red Hat Enterprise Virtualization Hypervisors.

Chris Curran

David Jorm

Stephen Gordon

Red Hat Enterprise Linux 5.7 Hypervisor Deployment Guide
The complete guide to installing, deploying and maintaining Red Hat Enterprise Virtualization Hypervisors.
Edition 6

Author	Chris Curran	*ccurran@redhat.com*
Author	David Jorm	*djorm@redhat.com*
Author	Stephen Gordon	*sgordon@redhat.com*

This book is a guide to installing and administering Red Hat Enterprise Virtualization Hypervisor version 5.7.

Preface

This book is a guide to installing Red Hat Enterprise Virtualization (RHEV) Hypervisors. This book provides a step-by-step procedures for installing Hypervisors with Red Hat Enterprise Virtualization Manager for Desktops. Alternative and advanced installation and configuration options are covered to assist users with varied environments and systems.

1. About this Guide

This guide describes how to install, setup and configure and Red Hat Enterprise Virtualization Hypervisors.

1.1. The Red Hat Enterprise Virtualization 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.

The Red Hat Enterprise Virtualization documentation suite

- *Red Hat Enterprise Virtualization Hypervisor Release Notes* contain release specific information for Red Hat Enterprise Virtualization Hypervisors.

- *Red Hat Enterprise Virtualization Manager Release Notes* contain release specific information for Red Hat Enterprise Virtualization Managers.

- *Red Hat Enterprise Virtualization for Desktops Installation Guide* describes the installation prerequisites and procedures. Read this if you need to install Red Hat Enterprise Virtualization for Desktops. The installation of hosts, manager and storage are covered in this guide. You will need to refer to the *Red Hat Enterprise Virtualization for Desktops Administration Guide* to configure the system before you can start using the platform.

- *Red Hat Enterprise Virtualization for Desktops Administration Guide* describes how to setup, configure and manage Red Hat Enterprise Virtualization for Desktops. It assumes that you have successfully installed the Red Hat Enterprise Virtualization manager and hosts.

- *Red Hat Enterprise Virtualization for Desktops User Guide* describes how users of the Red Hat Enterprise Virtualization for Desktops system can access and use virtual desktops.

- *Red Hat Enterprise Virtualization 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 for Desktops Installation Guide*

- *Red Hat Enterprise Virtualization Manager API Guide* describes how to use the command line functions to set up and manage virtualization tasks. Use this guide if you do not wish to use the graphical user interface, and prefer using the command line.

1.2. Audience

This documentation suite is intended for system administrators and OEMs installing Red Hat Enterprise Virtualization Hypervisors.

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*[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:

> 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).

---

[1] https://fedorahosted.org/liberation-fonts/

> 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 by email to the author of the manual, David Jorm (*djorm@redhat.com*). When submitting a bug report, be sure to mention the manual's identifier: *Deployment-Guide*.

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.

Introducing Red Hat Enterprise Virtualization Hypervisors

The Red Hat Enterprise Virtualization (RHEV) Hypervisor is a compact, full featured virtualization platform for quickly and easily deploying and managing virtualized guests. The RHEV Hypervisor is part of the Red Hat Enterprise Virtualization suite. The RHEV Hypervisor is designed for integration with the Red Hat Enterprise Virtualization Manager for Servers and the Red Hat Enterprise Virtualization Manager for Desktops.

The RHEV Hypervisor can be installed from USB storage devices, CD-ROMs, DVDs, pre-installed by an OEM or provisioned with PXE.

The RHEV Hypervisor is based on the Kernel-based Virtual Machine (KVM). KVM is an advanced and efficient virtualization hypervisor implemented as a Linux kernel module. As KVM is a kernel module, it leverages the existing Red Hat Enterprise Linux kernel and benefits from the default kernel's extensive testing, device support and flexibility.

The RHEV Hypervisor is constructed using a subset of packages from Red Hat Enterprise Linux along with several special RHEV packages. RHEV is certified for all hardware certified for Red Hat Enterprise Linux (except where noted, refer to *Chapter 3, Requirements and limitations of Red Hat Enterprise Virtualization Hypervisors*).

Understanding the Red Hat Enterprise Virtualization Hypervisor

This chapter provides a brief introduction to the technical aspects of the Red Hat Enterprise Virtualization Hypervisor. This chapter is written to assist technical users with identifying issues and understanding the Red Hat Enterprise Virtualization Hypervisor.

The Red Hat Enterprise Virtualization Hypervisor is derived from the Managed Node component of the oVirt project. While the core technology of oVirt Managed Node is used, the Red Hat Enterprise Virtualization Hypervisor incorporates its own management communications stack. For more information on the oVirt project, visit *oVirt.org*[1].

Red Hat Enterprise Virtualization Hypervisor directories

Red Hat Enterprise Virtualization Hypervisors have a directory layout that varies from Red Hat Enterprise Linux and other Linux Distributions. All Red Hat Enterprise Hypervisors have the following directories:

`/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 5MB.

`/boot`
> The **`/boot`** directory contains the boot loader, the kernel and the **`initramfs`** file.
>
> The **`/boot`** directory is not mounted by default. The **`ovirt-config-boot`** boot command mounts the partition only when needed for an upgrade.
>
> The **`/boot`** directory must be exactly 50MB.

`/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 livecd 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.

---

[1] http://ovirt.org

/var/vdsm/
> The **/var/vdsm/** is used by the vdsmd 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.

Requirements and limitations of Red Hat Enterprise Virtualization Hypervisors

This chapter contains all system requirements and limitations which apply to Red Hat Enterprise Virtualization (RHEV) Hypervisors. These requirements are determined by 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.

System requirements

Red Hat Enterprise Virtualization Hypervisors have the following hardware requirements:

- A CPU with the hardware virtualization extensions, this means an AMD system with AMD-V™ or an Intel system with Intel VT®.

  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.

  > ### ⭐ Virtualization extensions
  >
  > Some systems have the virtualization extensions disabled in BIOS; virtualization extensions must be enabled in BIOS. First, boot into rescue mode (by typing linux rescue at the RHEV Hypervisor boot prompt) or by entering the shell from the configuration menu. Determine if your processor has virtualization extensions with this command:
  >
  > ```
  > $ grep -E 'svm|vmx' /proc/cpuinfo
  > ```
  >
  > If any output is shown, the processor is hardware virtualization capable. However in some circumstances manufacturers disable the virtualization extensions in BIOS. 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.

- The CPU must have the Intel® 64 or AMD64 CPU extensions.

- At least one network controller with a minimum bandwidth of 1Gbps.

- At least 512MB of RAM.

- Sufficient RAM for virtual machines. The amount of RAM required for guests depends on:

- the guest operating system's requirements,

- the guests' application requirements, and

- memory activity and usage of guests.

  KVM can overcommit physical RAM for virtualized guests by only allocating RAM for guests as required and shifting underutilized guests into swap.

- The minimum supported internal storage for a Hypervisor is the total of the following list:

  - The root partitions require exactly 512MB of storage.

  - The boot partition requires exactly 50MB of storage.

  - The logging partition is required and the recommended size is 2048MB.

  - The default and minimum size for the configuration partition is 5MB.

  - Swap is required and the recommended size is the recommended swap partition (from this article on Red Hat *Knowledgebase*[1]) and the overcommit ratio multiplied by the RAM (the default overcommit ratio is 0.5). The minimum size for the swap partition is 5 MB.

    ```
    Recommended swap + (RAM * overcommit) = swap partition size
    ```

- The default size of the data partition is the size of the system's RAM plus 512MB. The minimum size is 256MB.

Example 3.1. Determining minimum storage requirements
This example covers the minimum required storage for a RHEV Hypervisor installation on a system with 14GB of RAM (approximately 14,336MB).

```
The root partitions          = 512MB
The boot partition          = 50MB
The logging partition       = 2048MB
The config partition        = 5MB
The swap partition (using a 0.6 overcommit ratio)
4GB (4096MB) + (0.6 * 14GB) = 12698MB
The data partition is the size of the system's RAM plus 512MB
             (14GB + 512MB) = 14750MB
-------------------------------------------------------------
                    Total = 30063MB
(
```

Note, that this is for the *minimum* storage for an installation. It is recommended to use the defaults which will use more storage space.

Limitations

The following limitations apply to hardware the Hypervisor is supported on and virtualized guests:

- A maximum of 64 physical CPUs on the host.

- A maximum of 1TB of RAM.

- A maximum of 16 virtualized CPUs per guest.

- A maximum of 256GB of virtualized RAM per 64 bit guest.

- A maximum of 4GB of virtualized RAM per 32 bit guest. Note, the guest may not register the entire 4GB. How much RAM the guest recognizes is limited by its operating system.

- A maximum of 8 virtualized storage devices per guest.

- A maximum of 8 virtualized network interface controllers per guest.

- A maximum of 32 virtualized PCI devices per guest.

RHEV virtualized guest support

RHEV 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)

- Windows XP Service Pack 3 and newer (32 bit only)

- 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)

- Windows 7 (32 bit and 64 bit)

Para-virtualized drivers 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.

Guest operating system	Para-virtualized drivers
Windows XP	Block and network drivers
Windows 7 (32 bit and 64 bit)	Block and network drivers
Windows Server 2008 (32 bit and 64 bit)	Block and network drivers
Windows Server 2003 R2 (32 bit and 64 bit)	Block and network 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

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 (RHEV) Hypervisor.

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

4.1. Preparation instructions

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

Entitlements to the `Red Hat Enterprise Virtualization Hypervisor 5` channel must be available to your RHN account to download the RHEV Hypervisor images.

Downloading and Installing the Package

The *rhev-hypervisor** package contains the Red Hat Enterprise Virtualization Hypervisor ISO image. The ISO itself contains additional tools for USB and PXE installations.

1. Download the latest *rhev-hypervisor** package from Red Hat Network onto a Red Hat Enterprise Linux system to which you have root access. The list of Hypervisor packages is available at *https://rhn.redhat.com/rhn/channels/PackageList.do?cid=9462*.

2. As `root`, navigate to the location of the downloaded package. Install it to the system in preparation for creation of the boot media.

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

3. Extract the **livecd-iso-to-pxeboot** and **livecd-iso-to-disk** commands from the ISO. To do this:

   a. Change to the directory the Hypervisor ISO is installed in, while logged in as `root`.

   ```
   # cd /usr/share/rhev-hypervisor
   ```

   b. Create a directory to which the ISO file will be mounted on a temporary basis.

   ```
   # mkdir iso/
   ```

   c. Mount the ISO file as a `loopback` device. Use the directory created in the previous step as the mount target.

   ```
   # mount -o loop rhev-hypervisor.iso iso/
   ```

d. Copy the **livecd-iso-to-pxeboot** and **livecd-iso-to-disk** scripts to the current directory.

```
# cp iso/LiveOS/livecd-iso-to-* ./
```

e. Unmount the ISO file and remove the temporary directory.

```
# umount iso/
# rmdir iso/
```

The RHEV Hypervisor ISO image is located in the **/usr/share/rhev-hypervisor/** directory and named **rhev-hypervisor.iso**.

BIOS settings and boot process troubleshooting

Before installing RHEV Hypervisors you should verify your BIOS is correctly configured for the installation method you intend to use. 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.

Not all motherboards support the boot methods described in this chapter. Consult your motherboard or system manufacturer's documentation on whether it is possible to use a particular boot method. That said, many modern systems support all of the boot methods listed in this chapter.

> ⚠️ **Warning**
>
> BIOS settings vary from manufacturer to manufacturer. Therefore, any examples of settings may be inaccurate for some systems. Due to this inconsistency, you should review your motherboard or system manufacturer's documentation.

Confirm hardware virtualization support

Verify that your system supports RHEV Hypervisors. RHEV Hypervisors require the virtualization extensions are present and enabled in the BIOS before installing the Hypervisor.

1. Boot the RHEV 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.

4.2. Deploying RHEV Hypervisors with PXE and tftp

This section covers installing RHEV 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 and extract the **livecd-iso-to-pxeboot** script. Refer to *Downloading and Installing the Package*.

2. Create **vmlinuz** and **initrd** images with **livecd-iso-to-pxeboot**:

```
# /usr/share/rhev-hypervisor/livecd-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.
```

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

    *   **pxelinux.0**

    *   **pxelinux.cfg/default**

    *   **vmlinuz0**

    *   **initrd0.img**

4. The **vmlinuz** and **initrd** files can be imported into a cobbler server or any other PXE and `tftp` servers.

    Import the files to the appropriate directory.

5. 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=/rhev-hypervisor.iso
 rootfstype=auto ro liveimg
              check crashkernel=128M@16M elevator=deadline processor.max_cstate=1
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.

> ⚠️ **Warning**
>
> The *root=/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 should be passed when calling **livecd-iso-to-pxeboot**. For example, for the ISO file **rhev_hypervisor_2_2.iso** use the command **livecd-iso-to-pxeboot rhev_hypervisor_2_2.iso**. This will produce the correct parameter *root=/rhev_hypervisor_2_2.iso* in **pxelinux.cfg/default**.

4.2.1. Booting a RHEV Hypervisor with PXE

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

To boot a RHEV 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 should begin.

5. Change or disable network booting after the RHEV Hypervisor is installed. This is to avoid overwriting the installation on each reboot (unless this is desired functionality) and to prevent certain security vulnerabilities.

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

> 💬 **Kernel parameters**
>
> For more information on the kernel parameters, refer to *Section 5.2, "RHEV Hypervisor Kernel Parameters and Automated Installation"*.

The RHEV Hypervisor should now be installed.

4.3. Preparing a RHEV Hypervisor USB storage device

RHEV Hypervisors can install from USB storage devices. They can also be installed to USB storage devices or 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.

> **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.

4.3.1. Making a bootable RHEV Hypervisor USB storage device

This section covers making bootable RHEV Hypervisors on USB storage devices.

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

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

```
# /usr/share/rhev-hypervisor/livecd-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 RHEV Hypervisor. The default RHEV Hypervisor image location is **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**. The **livecd-iso-to-disk** command requires devices to be formatted with the FAT or EXT3 file system.

> **Partitions and `livecd-iso-to-disk`**
>
> **livecd-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 and extract the **livecd-iso-to-disk** script. Refer to *Downloading and Installing the Package*.

2. Use the **livecd-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.

    ```
    # /usr/share/rhev-hypervisor/livecd-iso-to-disk --format --reset-mbr /usr/share/rhev-
    hypervisor/rhev-hypervisor.iso /dev/sdc
    Verifying image...
    /usr/share/rhev-hypervisor/rhev-hypervisor.iso:   eccc12a0530b9f22e5ba62b848922309
    ```

```
Fragment sums: 8688f5473e9c176a73f7a37499358557e6c397c9ce2dafb5eca5498fb586
Fragment count: 20
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/sdc1`) is ready to boot a RHEV Hypervisor.

4.3.2. Booting a RHEV Hypervisor USB storage device

Booting a RHEV 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 RHEV Hypervisor bootable USB storage device.

3. Restart the system

4. The RHEV Hypervisor should boot automatically.

If the Hypervisor is running, you must now initialize the local storage device. Refer to *Booting the Hypervisor for installation* for details.

4.4. Preparing a RHEV Hypervisor from a CD-ROM or DVD

RHEV Hypervisors can be installed with a CD-ROM or DVD.

4.4.1. Making a bootable RHEV Hypervisor CD-ROM

Burn the RHEV 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.

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

    If the package version is in the output, as seen above, the package is available.

    If it is not present, 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 RHEV Hypervisor image location, **/usr/share/rhev-hypervisor/rhev-hypervisor.iso**.

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

If no errors occurred, the RHEV 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 RHEV 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.

4.4.2. Booting a RHEV 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 RHEV Hypervisor CD-ROM and turn the computer on.

The RHEV Hypervisor should start to boot. If the RHEV 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 *Booting the Hypervisor for installation* for details.

Installing Red Hat Enterprise Virtualization Hypervisors

This chapter covers installing Red Hat Enterprise Virtualization (RHEV) Hypervisors.

Red Hat Enterprise Virtualization requires a RHEV Manager and one or more RHEV Hypervisors or Red Hat Enterprise Linux hosts. Installing the Red Hat Enterprise Virtualization Manager and Red Hat Enterprise Linux hosts is covered by the *Red Hat Enterprise Virtualization Manager for Servers Installation Guide*. This chapter covers installing Red Hat Enterprise Virtualization Hypervisors.

RHEV Hypervisors can use Storage Area Networks (SANs) and other network storage for storing virtualized guest images. RHEV Hypervisors can be installed on SANs, provided that the Host Bus Adaptor (HBA) can be configured as a boot device in BIOS.

RHEV Hypervisors can 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 can use multipath. This is most devices, except USB storage and some older ATA disks.

> **RHEV Manager installation**
>
> The RHEV Manager must be installed before RHEV Hypervisors are installed. Refer to the *Red Hat Enterprise Virtualization for Servers Installation Guide* for instructions on installing the RHEV Manager.

There are two methods for installing RHEV Hypervisors:

* Installing with the Hypervisor Configuration Menu menu (refer to *Section 5.1.1, "The RHEV Hypervisor Configuration Menu"*

* Installing with kernel command line parameters (refer to *Section 5.2, "RHEV Hypervisor Kernel Parameters and Automated Installation"*

5.1. Red Hat Enterprise Virtualization Hypervisor interactive installation

RHEV Hypervisors must be installed on a physical servers, not virtual machines.

The instructions in this section can be used for installing on multiple systems. Always remember to use unique hostnames and IP addresses to avoid networking conflicts.

Booting the Hypervisor for installation

There are several methods for booting Hypervisors, refer to *Chapter 4, Preparing Red Hat Enterprise Virtualization Hypervisor installation media* for detailed instructions on preparing boot media for RHEV Hypervisor installation.

Procedure 5.1. Booting the Hypervisor
1. Insert the RHEV Hypervisor installation media.

2. Power on the system and ensure the system boots from the installation media.

3. Once the RHEV Hypervisor has completed the boot phase, the kernel parameter prompt appears.

```
ISOLINUX 3.11 2005-09-02  Copyright (C) 1994-2005 H. Peter Anvin
This is the Red Hat Enterprise Virtualization Hypervisor 5.7-2.2 (4.2)
- to boot press the <ENTER> key
or type linux and press the <ENTER> key.
- Extra kernel command line parameters can be passed following the
linux keyword.
- Boot parameters summary:
BOOTIF=link|eth*|<MAC>(pxelinux format)
storage_init=[usb|scsi|cciss[:serial#]]|/dev/...
storage_vol=:SWAP_MB::CONFIG_MB:LOGGING_MB:DATA_MB
mem_overcommit=<overcommit_ratio>
upgrade
ip=<client-ip> netmask=<netmask> gateway=<gw-ip>
vlan=<VLAN-id>
dns=server[:server]
ntp=server[:server]
hostname=fqdn
management_server=RHEVM-host[:port]
netconsole_server=RHEVM-host[:port]
boot:
```

Enter any required kernel arguments. For example:

```
:linux firstboot
```

Kernel parameters

For more information on the kernel parameters, refer to *Section 5.2, "RHEV Hypervisor Kernel Parameters and Automated Installation"*.

Upgrading existing hypervisors

To upgrade an existing hypervisor, you must boot the kernel with the *upgrade* parameter. This will automatically upgrade and reboot the system, rather than taking you to the interactive configuration menu. For more information, refer to *upgrade*.

PXE booting

You can append the kernel boot arguments to the PXE configuration file (**/pxelinux.cfg/ default**). This can be used to run an automated setup, covered in *Section 5.2, "RHEV Hypervisor Kernel Parameters and Automated Installation"*, and may be more useful depending on your environment.

5.1.1. The RHEV Hypervisor Configuration Menu

Once the RHEV Hypervisor boots you will see the Hypervisor Configuration Menu.

The Hypervisor Configuration Menu is an interactive tool for installing Red Hat Enterprise Virtualization Hypervisors. The Hypervisor Configuration Menu resembles *Example 5.1, "Red Hat Enterprise Virtualization Hypervisor Configuration Menu"*. Throughout this chapter it will be referred to as the main menu.

Example 5.1. Red Hat Enterprise Virtualization Hypervisor Configuration Menu

```
Red Hat Enterprise Virtualization Hypervisor release 5.7-2.2

 Hypervisor Configuration Menu

1) Configure storage partitions    6) Configure the host for RHEV
2) Configure authentication         7) View logs
3) Set the hostname                 8) Install locally and reboot
4) Networking setup                 9) Support Menu
5) Register Host to RHN
Choose an option to configure:
```

Configuration menu options

- **Configure storage partitions**: prepares the local storage device for installation. Refer to *Section 5.1.3, "Partitioning the disk"*.

- **Configure authentication**: sets the root password for administration and SSH access. Refer to *Section 5.1.4, "Set the administrator password"*.

- **Set the hostname**: sets the hostname manually instead of using DHCP. Refer to *Section 5.1.5, "Set the hostname"*.

- **Networking setup**: configures network settings and services including: DHCP, IPv4, NTP and DNS. Refer to *Section 5.1.6, "Setup Networking"*.

- **Register Host to RHN**: registers the host to the Red Hat Network (RHN). Refer to *Section 5.1.7, "Register to RHN"*.

- **Configure the host for RHEV**: enter IP address of the RHEV Manager server. Refer to *Section 5.1.8, "Configure the host for RHEV"*.

- **View logs**: view log files to debug the installation procedure. Refer to *Section 5.1.9, "View logs"*.

- **Install locally and reboot**: install the Hypervisor to the storage device and reboot. Refer to *Section 5.1.10, "Installation"*.

- **Support Menu**: contains options for continuing stateless boot and exiting to a shell prompt. This menu is not supported except when working with Red Hat support.

5.1.2. Remove existing RHEV Hypervisor installations

If the Hypervisor was previously installed on the local system, remove the existing installation by selecting the **Support Menu** option from the main menu, then selecting **Uninstall an existing Hypervisor**.

> ## On systems with existing Hypervisors
>
> On systems that have a RHEV Hypervisor installed and some systems with LVM volumes, the *firstboot* parameter may be required to trigger the interactive installation. To use the *firstboot* parameter, append it to other kernel parameters. On a system with an existing RHEV Hypervisor use the following kernel parameters to activate the interactive installation:
>
> ```
> : linux firstboot
> ```

For more information on uninstalling previous versions of the RHEV Hypervisor refer to *Appendix A, Uninstalling Red Hat Enterprise Virtualization Hypervisors from local storage devices*.

5.1.3. Partitioning the disk

The partitions used by the RHEV Hypervisor are configured from this menu. For most installations the default partition layout is optimal.

> ## Warning
>
> Running this procedure will destroy all data on the local storage device selected with the *storage_init* parameter.

Procedure 5.2. The Storage Configuration menu

1. Select the **Configure storage partitions** option from the main menu to open the **Configure storage partitions** menu.

    Example 5.2. The **Configure storage partitions** menu

    ```
     Configure storage partitions

    1) Configure
    2) Review
    3) Commit configuration
    4) Return to the Hypervisor Configuration Menu
    Choose an option:
    ```

2. **Select storage device**
    Select **Configure**.

    If you have more than one storage device available, you are prompted to select the storage device on which to install the RHEV Hypervisor. If there is only one storage device, skip to *Step 3*.

    Example 5.3. Select storage device

    ```
    /dev/mapper/SServeRA_venh_076A0444 (sda) (69889 MB)
    Disk Identifier: storage_serial_SServeRA_venh_076A0444
    /dev/sdb (15296 MB)
    Disk Identifier: storage_serial_UFD_USB_Flash_Drive_XFHO0KCK69T746WFHAUB
    1) /dev/mapper/SServeRA_venh_076A0444  3) Abort
    2) /dev/sdb                            4) Manual selection
    ```

Select the appropriate storage device from the menu. If the device you wish to use is not shown, select **Manual selection**. You will then be prompted to enter the path to the storage device. The path must be entered exactly; wildcards are not permitted.

3. **Choose partition sizes**
   You are prompted to set the size of boot, swap, root, config, logging and data partitions. The value is accepted in megabytes (MB).

   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.

   Pressing `Enter` without any value sets the partition to the default size.

   > ⭐ **Important**
   >
   > Do not append units onto the end of the partition size values. Using MB or similar unit on the end of a value may have adverse effects. For example, use the value **5** instead of **5MB**.

   a. The boot partition contains the boot loader, kernel and initramfs.

      The default and only supported size for the boot partition is 50MB. The size of the boot partition cannot be configured.

   b. The swap partition is used for swapping underused memory to the hard drive to speed up memory performance. The default size of the swap partition is calculated from amount of RAM and overcommit ratio (default is 0.5). RHEV 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 5MB.

      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 overcommit ratio (RAM multiplied by the overcommit ratio).

      ```
      Recommended swap + (RAM * overcommit) = swap partition size
      ```

      Pressing `Enter` during the installation sequence sets a recommended value for the swap partition.

   c. The root partition contains the core files for the Hypervisor and the `LiveOS` directory from the installation media. The default total size of the root file system partitions is 512MB. That is broken up into 256MB for one root partition and an additional 256MB for a backup root partition.

      The default and only supported size for the root partition is 256MB. The size of the root partition cannot be configured.

   d. The config partition stores configuration files for the Hypervisor.

      The default and minimum size for the configuration partition is 5MB.

   e. The logging partition stores all logs for the Hypervisor.

The logging partition is required and the recommended size is 2048MB.

f. The data partition stores core files and kdump files and acts as temporary storage for ISO files. The default size of the data partition is the remaining available disk space (labeled as -1).

When ISO files are uploaded using the ISO Uploader tool, they are temporarily stored on the data partition of a RHEV Hypervisor host, before being moved to the shared ISO Library. The data partition must include sufficient space to store these ISO files, which can be over 4GB in size for modern operating systems.

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 kdumps. A kdump file is usually the same size the host's system RAM. The data partition also stores the RHEV Hypervisor ISO file for RHEV Hypervisor upgrades.

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

4. **Review**

Select the **Review** menu option. The review menu output is similar to the screenshot below:

```
The local disk will be repartitioned as follows:
===============================================
             Physical Hard Disk: /dev/hdb (10000 MB)
Disk Identifier: storage_serial_QM00002
           Boot partition size: 50 MB
           Swap partition size: 2233 MB
  Installation partition size: 256 * 2 MB
 Configuration partition size: 5 MB
        Logging partition size: 2048 MB
           Data partition size: 5152 MB
```

5. Select **Commit configuration** to accept the partitions. The following warning displays:

6.
```
!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!
!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!
!!WARNING!!                                                   !!WARNING!!
!!WARNING!!                                                   !!WARNING!!
!!WARNING!!    If you proceed, all data on your selected storage  !!WARNING!!
!!WARNING!!    device will be destroyed and your hard disk        !!WARNING!!
!!WARNING!!    will be irreversibly reconfigured.                 !!WARNING!!
!!WARNING!!                                                   !!WARNING!!
!!WARNING!!                                                   !!WARNING!!
!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!
!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!WARNING!!

Do you wish to proceed([Y]es or [N]o)?
```

7. Press **Y** to proceed and format the partitions.

Partitions are created and file systems are formatted. Once complete, the main menu reappears.

5.1.4. Set the administrator password

Setting a strong and easy to remember administrator password is an essential part of securing your infrastructure.

Setting an administrator password is not required if you do not need SSH remote access. SSH remote access is not supported unless you are instructed to enable it by Red Hat support. In normal usage scenarios you should not need to set a root password.

Procedure 5.3. Set the administrator password

1. Select **Configure authentication** from the main configuration menu.

    > Example 5.4. The **Configure authentication** menu
    >
    > ```
    >  SSH remote access is currently disabled
    >
    > 1) Set administrator password
    > 2) Toggle SSH password authentication
    > 3) Return to the Hypervisor Configuration Menu
    > Choose an option to configure:
    > ```

2. Select **Set the administrator password**.

3. The password prompt displays.

    ```
    Set the system administrator's (root) password:
    New UNIX password:
    ```

    Enter the preferred password. You will be prompted to retype the password.

4.
    ```
    Changing password for user root.
    passwd: all authentication tokens updated successfully.
    storing to /config :
    /etc/shadow File persisted
    ```

    This message indicates that the password has been set and saved to local storage.

Configure SSH password authentication

SSH password authentication is required for accessing RHEV Hypervisors with the SSH protocol. By default, SSH password authentication is disabled. SSH remote access is not supported unless you are instructed to enable it by Red Hat support.

To enable SSH password authentication:

1. Select **Toggle SSH password authentication** from the **Configure authentication** menu.

2.
    ```
    Enable SSH password authentication ([Y]es or [N]o)?
    ```

    Press **y** to enable SSH password authentication.

3. If successful, the following message is displayed:

    ```
    SSH remote access is currently enabled.
    ```

Select **Return to the Hypervisor Configuration Menu** to continue the installation.

5.1.5. Set the hostname

The Red Hat Enterprise Virtualization Manager requires a fully resolvable hostname for each RHEV Hypervisor. If you are using DHCP and DNS a hostname can be assigned to the Hypervisor at boot time.

If the DHCP server does not set hostnames, the hostname must be specified manually (with the process below). This hostname must resolve to the assigned IP address located with a DNS lookup request.

1. Select the **Set the hostname** option from the main menu.

2. Enter the hostname in the following format: **anyname.domain.example.com**. Where **anyname** is the individual name for the host on the specified domain.

   ```
   What is this Hypervisor's hostname?
   ```

   Enter the hostname for the Hypervisor and press **Enter**.

3. If the hostname setting procedure is successful, this message is displayed:

   ```
   The hostname is set.
   ```

You are returned to the Configuration menu.

Resetting hostnames

To reset a hostname:

1. Open the **Set the hostname** menu.

2. The specify a hostname prompt appears:

   ```
   What is this Hypervisor's hostname?
   ```

   Press the Enter key without typing anything else to reset the hostname.

The hostname is now reset to use DHCP.

5.1.6. Setup Networking

Configure your system's network settings.

> ⭐ **Multiple networks**
>
> For installation purposes only a single network can be configured for access to the RHEV Manger. Once the RHEV Hypervisor is installed, additional network interfaces can be configured with the RHEV Manager user interface.

1. Select the **Networking setup** option.

2. The network devices attached to your system are listed.

3.
   ```
   Choose an interface or a configuration option:
   ```

Select the network interface which will be used to connect to the RHEV Manager. On systems with one network interface there is only `eth0` to choose from.

> ⭐ **Device support**
>
> Wireless and logical network devices are presently unsupported.

4. **Identifying the interface**

   Hypervisors include a tool to assist in physically identifying which network interface was selected.

   ```
   Help identify eth0 by blinking the lights for 10 seconds ([Y]es or [N]o)?
   ```

   Press **y** to identify which socket *eth0* is. This is useful for systems with multiple network interfaces.

   If assistance is not required, continue by pressing **n**.

   If the selected network interface does not support this feature, the system will display the following message:

   ```
   Blinking is not supported by your network interface.
   ```

5. **VLANs**

   A prompt will be displayed asking whether to include VLAN support:

   ```
   Include VLAN support ([Y]es, [N]o or [A]bort)?
   ```

   If VLAN support is not required, press **n**.

   To enable VLAN support, enter **y**.

   - ```
     What is the VLAN ID for this device (a=abort)
     ```

     Enter the VLAN ID number when prompted.

6. **IPv4**

   The IPv4 prompt appears.

   ```
   Enable IPv4 support ([S]tatic IP, [D]HCP, [N]o or [A]bort)?
   ```

   Select **DHCP** or **Static IP** depending on your networking environment.

   To use DHCP, press **D** to enable IPv4 support with DHCP. It is required to statically map all IP addresses for each RHEV Hypervisor on your DHCP server.

   Alternatively, assign a static IP address by pressing **S**. You are then prompted for the IP address, Netmask and gateway.

   Confirm your settings.

```
Is this correct ([Y]es, [N]o or [A]bort)?
```

Press **Y** if correct to continue.

7. **DNS**

   Configure DNS if using static network configuration or of your DHCP server is not assigning DNS servers.

   a. Enter the IP address of the first DNS gateway or press **Enter** to abort.

   ```
   Enter the first DNS server (or press the Enter key to exit):
   ```

   b. Enter the IP address of the second DNS gateway or press **Enter** to skip.

   ```
   Enter the second DNS server (or press the Enter key to exit):
   ```

   c. Confirm your settings.

   ```
   Is this correct ([Y]es, [N]o or [A]bort)?
   ```

   Press **Y** if correct to continue.

8. **Network Time Protocol (NTP)**

   The Network Time Protocol (NTP) keeps system clocks in time with the rest of the world. NTP is crucial for some host systems with inconsistent clocks. Inaccurate timekeeping prevents Hypervisors from migrating virtual machines.

   a.
   ```
   Enter an NTP server:
   ```

   Enter the IP address or hostname of an NTP server or press **Enter** to abort.

   b. You will be prompted to enter additional NTP servers. Enter as many as required and press **Enter** on a blank line to save your selection.

9. Select **Save And Return To Menu** to save the configuration.

10. If the network is configured correctly, the output should be similar to *Figure 5.1, "Working network configuration"*.

Figure 5.1. Working network configuration

You are now returned to the main menu.

5.1.7. Register to RHN

This procedure registers the host to RHN. Registering the host to RHN allows you to apply Red Hat Enterprise Linux virtualization or virtualization platform entitlements to the Hypervisor host. If you are not using these entitlements, you do not need to register the host to RHN.

- Select **Register Host to RHN** from the main menu.

    To register to the public RHN system, select **1) Register to public RHN**

    a.
    ```
    Enter RHN account username:
    ```

    Enter your RHN account username.

    b.
    ```
    Enter password for RHN account:
    ```

    Enter your RHN account password.

    c.
    ```
    Enter profile name for this system (optional):
    ```

    Enter a profile name for this system. This parameter is optional. If a blank profile name is provided, the system hostname will be used by default.

    d.
    ```
    Enter HTTP proxy in host:port format, e.g. proxy.example.com:3128 (leave blank
      if not required):
    ```

    Enter your HTTP proxy address. This parameter is optional. If a blank proxy address is provided, no proxy server will be used.

    To register to an RHN Satellite, select **2) Register to RHN Satellite**

    a.
    ```
    Enter URL of RHN Satellite (ie. https://satellite.example.com)
    ```

    Enter your the URL of your RHN Satellite.

    b.
    ```
    Enter location of satellite cert (https://satellite.example.com/pub/RHN-ORG-
    TRUSTED-SSL-CERT):
    ```

    Enter the URL of your satellite certificate, or press enter to select the default URL shown.

    c.
    ```
    Enter RHN account username:
    ```

    Enter your RHN account username.

    d.
    ```
    Enter password for RHN account:
    ```

    Enter your RHN account password.

    e.
    ```
    Enter profile name for this system (optional):
    ```

    Enter a profile name for this system. This parameter is optional. If a blank profile name is provided, the system hostname will be used by default.

f.
```
Enter HTTP proxy in host:port format, e.g. proxy.example.com:3128 (leave blank
   if not required):
```

Enter your HTTP proxy address. This parameter is optional. If a blank proxy address is provided, no proxy server will be used.

5.1.8. Configure the host for RHEV

This procedure connects the RHEV Hypervisor to a RHEV Manager.

When specifying the address of the RHEV Manager, use the following format for IP addresses or hostnames.

> ⭐ **Before you continue...**
>
> Ensure that you have configured the local storage and networking correctly before proceeding. If something is not correct change it now because you cannot change it after this point. To change a setting after this point you must repeat the procedure from the beginning.
>
> Ensure that the IP address or host name specified is fully resolvable.

Table 5.1. Red Hat Enterprise Virtualization Manager addressing examples

Type	Format	Example
IP address	IP[:port]	127.0.0.1:443
Host Name	host.name.domain[:port]	www.redhat.com:443

1. Select **Configure the host for RHEV** from the main menu.

2. **Connecting RHEV Manager**
   Once the **Configure the host for RHEV** menu item is selected the following output displays:

   ```
   Enter the RHEV Manager's hostname or IP address.
   Optionally: append a port after the hostname or IP address
   For example, 10.0.0.1:443 or rhev.example.com:443
   ```

   a. Enter the RHEV Manager's hostname or IP address in the specified format.

      If the RHEV Manager port is skipped the setup will use the default port, port 443. However, if a port has previously been assigned, the RHEV Hypervisor uses that port as the default instead. For example, if port 12345 was specified during a previous installation then the default port for future installations is port 12345.

      > 💬 **Using other ports**
      >
      > RHEV Managers use 443 by default. If you configured another port when installing the RHEV Manager write it down. Once another port is specified for the RHEV Hypervisor it will be recorded as the new "default".

   b. If the RHEV Manager is successfully located, the following output is displayed:

```
The RHEV Manager's address is set
The RHEV Manager's port is set.
```

3. **NetConsole access**

```
Enter the NetConsole manager's hostname or IP address.
Optionally: append a port after the hostname or IP address
For example, 10.0.0.1:25285 or rhev.redhat.com:25285
If you are unsure, use the RHEV Manager's address, without a port.
```

   a. Enter the NetConsole server and port. In most cases, this server is the RHEV Manager. The
       NetConsole port is the NetConsole port number specified during the installation of the RHEV
       Manager.

   b. If the NetConsole server is successfully found, the following output is displayed before being
       returned to the configuration menu:

```
The NetConsole manager address is set.
The NetConsole manager port is set.
```

5.1.9. View logs

Various log files can viewed to assist in debugging the installation procedure. Once a file is selected by
its menu option number, it is displayed in full on the screen. You can scroll through the file contents by
pressing **Shift**+**Page Up** to scroll up and **Shift**+**Page Down** to scroll down. The log files available
are:

/var/log/ovirt.log

Contains detailed information about the hypervisor installation and configuration procedure.

/var/log/messages

Contains general messages from the system and core services.

/var/log/secure

Contains authentication and security messages, including failed login attempts.

5.1.10. Installation

Once all the previous steps are complete, the RHEV Hypervisor can be installed to the local storage
device.

Storage, networking and the RHEV Manager *must* be configured before continuing.

1. Select **Install locally and reboot** from the Hypervisor Configuration Menu.

2.
```
This option exits the menu and requires a reboot.
Ensure local storage and networking is configured
correctly on the previous menu before continuing.

The current kernel boot parameters are
crashkernel=128M@16M elevator=deadline processor.max_cstate=1
```

```
Enter additional kernel boot parameters (press Enter to use the present arguments):
```

Press **Enter** to continue.

3.
```
Do you wish to proceed ([Y]es or [N]o)?
```

Press **Y** to continue.

The final installation sequence begins. Once the installation is complete the RHEV Hypervisor will reboot.

After rebooting, the RHEV Hypervisor is installed and integrated with Red Hat Enterprise Virtualization Manager for Servers. You should have a new host available in the Red Hat Enterprise Virtualization Manager interface.

> **Note**
>
> You may need to remove the boot media and change the boot device order to prevent the installation sequence restarting after the system reboots.

5.2. RHEV Hypervisor Kernel Parameters and Automated Installation

This section covers the kernel command line parameters for RHEV Hypervisors. These parameters can be used to automate installation. The parameters are described in detail and an example is provided for an automated installation.

This installation method is an alternative to the interactive installation covered by *Section 5.1.1, "The RHEV Hypervisor Configuration Menu"*. Using the method covered in this chapter with a PXE server can, with some configuration, deploy multiple RHEV 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.

How the kernel arguments work

Below is a description of the RHEV 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 RHEV scripts and start up procedures.

3. An automated installation begins if all the required parameters are set.

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.

5.2.1. Storage parameters

The following parameters configure local storage devices for installing a RHEV Hypervisor.

storage_init

The *storage_init* parameter is required for an automated installation.

The *storage_init* parameter initializes a local storage device. Presently, RHEV Hypervisors use one storage device for local installation. There are four methods for defining which disk to initialize and install on.

a. For USB storage devices, use the *usb* parameter to select disk type, as reported by **udev ID_BUS**. Optionally, append a serial number to explicitly define which device to use. For example:

```
storage_init=usb:Generic_STORAGE_DEVICE_0000145418-0:0
```

b. For SCSI hard drives, use **storage_init=scsi:***serialnumber*. Optionally, append a serial number to explicitly define which device to use. For example:

```
storage_init=scsi:SCSI_HITACHI_HTS5416_SBADDFF
```

c. For CCISS devices, use **storage_init=cciss:***serialnumber*. Optionally, append a serial number to explicitly define which device to use. For example:

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

d. 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
```

Note that these device names are not guaranteed the same every time the system boots.

> ### Finding udev serial numbers
>
> Serial numbers can be found with this script. The script lists serial numbers for all disks attached to the system.
>
> ```
> for d in /dev/sd?; do eval $(udevinfo -q env -n $d); echo $d $ID_SERIAL;done
> ```

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 values separated by colons.

Values can be left undefined. If a value is not defined the installation will use the default value. 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 5.5. Format of the storage_vol parameter**
>
> ```
> storage_vol=:SWAP::CONFIG:LOGGING:DATA
> ```

- *SWAP*

  The swap partition is used for swapping underused memory to the hard drive to speed up memory performance. The default size of the swap partition is calculated from amount of RAM and overcommit ratio (default is 0.5). RHEV 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 5MB.

  Red Hat *Knowledgebase*[2] has an article on determining the size of the swap partition.

  Use the formula from the Red Hat Knowledgebase and add storage for the overcommit ratio (RAM multiplied by the overcommit ratio).

  ```
  Recommended swap + (RAM * overcommit) = swap partition size
  ```

  Leaving the value empty sets a 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 5MB.

- *LOGGING*

  The logging partition stores all logs for the Hypervisor. The logging partition is required and the recommended size is 2048MB.

- *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 kdumps. A kdump file is usually the same size the host's system RAM. The data partition also stores the RHEV Hypervisor ISO file for RHEV Hypervisor upgrades.

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

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

Example 5.6. Example: Using the storage_vol parameter to partition default sizes

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

Example 5.7. Example: Using the storage_vol parameter to partition certain sizes sizes

The Boot partition is defined as the fixed size of 50MB.

The Swap partition is defined as 4000MB.

The Root partition is defined as the fixed size of 512MB.

The Config partition is defined as 5MB.

The Logging partition is defined as 2048MB.

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

```
storage_vol=:4000::5:2048:-1
```

5.2.2. Networking and RHEV Manager 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 RHEV 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 RHEV Manager.

> ⭐ **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 you use the IPAPPEND 2 directive you do not need 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 into RHEV 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 5.8. 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 5.9. Using an IP address**
>
> ```
> ip=192.168.1.1
> ```

> **Example 5.10. Using DHCP**
>
> ```
> ip=dhcp
> ```

netmask

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

> **Example 5.11. Using a subnet mask**
>
> ```
> netmask=255.255.255.0
> ```

gateway

The *gateway* parameter sets the Internet gateway.

> **Example 5.12. 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 5.13. 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 5.14. 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 5.15. 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 RHEV Manager. Set this parameter if you are using VLANs.

> **Example 5.16. Using a VLAN**
>
> ```
> vlan=VLAN-ID
> ```

5.2.3. 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. If you are registering to a satellite server, the *rhn_url* parameter must be provided.

rhn_username

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

> **Example 5.17. Setting a username of testuser**
>
> ```
> rhn_username=testuser
> ```

rhn_password

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

> **Example 5.18. Setting a password of testpassword**
>
> ```
> rhn_password=testpassword
> ```

rhn_activationkey

The *rhn_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 *rhn_activationkey* and *rhn_username* are provided, the *rhn_activationkey* value will be used.

> **Example 5.19. Setting an activation key**
>
> ```
> rhn_activationkey=7202f3b7d218cf59b764f9f6e9fa281b
> ```

rhn_url

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

> **Example 5.20. Setting a satellite server URL**
>
> ```
> rhn_url=https://your-satellite.example.com
> ```

rhn_ca_cert

The *rhn_ca_cert* parameter sets the URL of the CA certificate used to connect to the satellite server. If it is not provided, the default value is *rhn_url*/pub/RHN-ORG-TRUSTED-SSL-CERT

> **Example 5.21. Setting a CA certificate URL**
>
> ```
> rhn_ca_cert=https://your-satellite.example.com/pub/RHN-ORG-TRUSTED-SSL-CERT
> ```

rhn_profile

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

> **Example 5.22. Setting a profile name of testhost**
>
> ```
> rhn_profile=testhost
> ```

5.2.4. Other parameters

management_server

The *management_server* parameter is required for an automated installation.

The *management_server* parameter sets the address of the RHEV 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 5.23. Connecting to a RHEV Manager at rhev.example.com:443**
>
> ```
> management_server=rhev.example.com:443
> ```

netconsole_server

The *netconsole_server* parameter sets the address of the RHEV Manager. The *netconsole_server* parameter takes an IP address or fully qualified domain name and, optionally, a port (the default port is 25285).

> **Example 5.24. Setting a NetConsole server at rhev.example.com:25285**
>
> ```
> netconsole_server=rhev.example.com:25285
> ```

mem_overcommit

The *mem_overcommit* parameter specifies the multiplier to use for adding extra swap to support memory overcommit. The default overcommit value is 0.5.

> **Example 5.25. Changing the memory overcommit to 0.7**
>
> ```
> mem_overcommit=0.7
> ```

firstboot

The *firstboot* parameter starts the interactive configuration menu. On systems that have a RHEV Hypervisor installed and some systems with LVM volumes, the *firstboot* parameter may be required to trigger the interactive installation. Refer to *Section 5.1.1, "The RHEV Hypervisor Configuration Menu"* for installation from the configuration menu.

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.

local_boot

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

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.

qemu_pxe

The *qemu_pxe* parameter is used to select which network bootloader is used in virtual machines. Possible values are *gpxe* and *etherboot*. For compataibility with RHEV Hypervisor 5.4-2.1, the default value is *etherboot*.

> **Example 5.26. Using the gpxe bootloader**
>
> ```
> qemu_pxe=gpxe
> ```

5.2.5. Example: Automated RHEV Hypervisor Installation

This example uses the kernel command line parameters for an automated RHEV Hypervisor installation. Ensure you customize the example for your environment as not all installations are the same.

> ⚠️ **Warning**
>
> This example may not work accurately on all systems. Review the descriptions of the parameters above and modify as required for your systems.

The following is a typical example for installing a RHEV Hypervisor with the kernel command line parameters

In this example, the RHEV Manager is located at the hostname: **rhevm.example.com**.

```
:linux storage_init=/dev/sda storage_vol=::::: local_boot BOOTIF=eth0
 management_server=rhevm.example.com netconsole=rhevm.example.com
```

> 💬 **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.

Maintaining and Administrating Red Hat Enterprise Virtualization Hypervisors

This chapter covers administration and maintenance tasks for Red Hat Enterprise Virtualization Hypervisors.

6.1. Enabling kdumps

RHEV 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.

All kdump files are placed on the **/data** partition of the storage on the RHEV Hypervisor host. When installing a RHEV Hypervisor host, ensure there is sufficient persistent storage space for kdump files. The size of the kdump file is proportional to the amount of RAM installed on the host. For a system with 32GB of RAM, a Hypervisor requires a **/data** partition of at least 32GB plus additional space for RHEV specific data (approximately 512MB).

The kdump file can be transferred off the RHEV Hypervisor host once the host is rebooted and is running properly. To transfer the kdump file, use the following procedure.

> ### SSH for Windows systems
>
> There is an Open Source SSH client for Windows systems called PuTTY. Download the client from the PuTTY homepage: *http://www.chiark.greenend.org.uk/~sgtatham/putty/*. Using PuTTY to transfer files is not supported by this guide although many of the steps will remain the same.

1. Access a Linux, UNIX, or similar operating systems with SSH and the **scp** command installed. The **scp** command has the following format:

   ```
   scp source-user@source_IP_OR_hostname:/path/to/files destination-
   user@destination_IP_OR_hostname:/path/to/destination
   ```

2. This example copies all kdump files to the **/tmp/RHEVkdumps** directory on the local system. Use the administrator (root) username and password configured during RHEV Hypervisor installation. Replace the example hostname, *rhevh.example.com*, with a valid IP address or hostname for the RHEV Hypervisor.

   ```
   $ scp -r root@rhevh.example.com:'/data/core/*' /tmp/RHEVkdumps
   ```

3. Once this procedure completes you can delete all the original kdump files on the Hypervisor. Log into the Hypervisor as root and execute the following.

   ```
   # rm -rf /data/core/*
   ```

The kdump file is now removed and available on the other host for support and debugging.

6.2. Setting the root password

By default, the root login and password are disabled on RHEV Hypervisors for security reasons. During unattended installs, there is no way to override this and set the root password at install time. In some circumstances, however, having root access to the RHEV Hypervisor host is necessary for debugging and support. To work around this, administrators must first gain temporary root access to the RHEV Hypervisor host, then set the password.

1. **Gaining temporary root access**

    Method 1: SSH using a key pair. This procedure is run from the RHEV Manager server.

    a. Download the latest versions of `putty.exe` and `puttygen.exe` to your RHEV Manager server from *http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html*

    b. Use puttygen to generate a new private key.

        i.   Run `puttygen.exe`. From the top menu, select **conversion** -> **Import key**

        ii. Select `C:\Program Files\RedHat\RHEVManager\Service\ca\keys\rhevm.pem`

        iii. Enter the passphrase for this key: mypass

        iv. Click **save private key**

    c. Use the private key to connect to the RHEV Hypervisor host.

        i.   Run `putty.exe`. Enter the IP address of the RHEV Hypervisor host you wish to connect to.

        ii. From the left panel select **SSH** -> **Auth**

        iii. Browse to the private key you just generated, click **open**

        iv. Enter the username root and use the password mypass. Open the SSH connection. You should now have a root shell open on the RHEV Hypervisor host.

    Method 2: Boot into rescue mode. This procedure requires physical access to the RHEV Hypervisor host, and is run from its physical console.

    a. Reboot the machine and wait for the GRUB prompt. Hit `enter` to interrupt normal booting.

    b. Modify the kernel line to boot into rescue mode. To do this, append the string `rescue` to the kernel line, separated from the rest of the line with a `space` character. Once the system boots, you will be dropped to a a single user root shell. The system must be rebooted into multi-user mode after following the instructions below.

2. **Setting the root password**

    Once you have successfully gained root access, you can set the root password for future use.

    a. Run the command `setup`

    b. Select option #1 **Configure authentication**

    c. Select option #2 **Toggle SSH password authentication**

    d. Enable SSH password authentication

e. Select option #1 **Set administrator password**

f. Enter your new root password

g. Select option #3 **Return to the Hypervisor Configuration Menu**

h. Select option #7 **Exit Hypervisor Configuration Menu**

Upgrading Red Hat Enterprise Virtualization Hypervisors

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

7.1. Upgrading a RHEV Hypervisor with a RHEV Manager

You can upgrade and reinstall a Red Hat Enterprise Virtualization Hypervisor host from an ISO stored on the RHEV 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**
>
> Upgrading hosts involves shutting down, deactivating guests and restarting the physical server. If any virtual servers 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 RHEV Hypervisor RPM package from *Red Hat Network*[1].

   Or, download the latest RHEV Hypervisor ISO image from *Red Hat Network*[2]

2. Copy the ISO image file to the `C:\Program Files\RedHat\RHEVManager\Service\RHEV-H Installer` folder on the Red Hat Enterprise Virtualization Manager server.

3. Change the file permissions so that `NETWORK SERVICE` has full control of the iso file.

   1. Use Windows Explorer to navigate to `C:\Program Files\RedHat\RHEVManager\Service\`.

   2. Right-click on the updated file, for example, `rhev-hypervisor.iso`, and select the Properties option. The **Properties** dialog displays.

   3. Click the **Security** tab, and select **NETWORK SERVICE** from the Group or user names panel.

   4. In the **Permissions** panel, click the **Allow Full Control** check box.

   5. Click **OK**.

> ![Important icon] **Important**
>
> 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 7.1. Upgrading a RHEV 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, perform a search.

2. On the **Details** pane, click the **General** Tab.
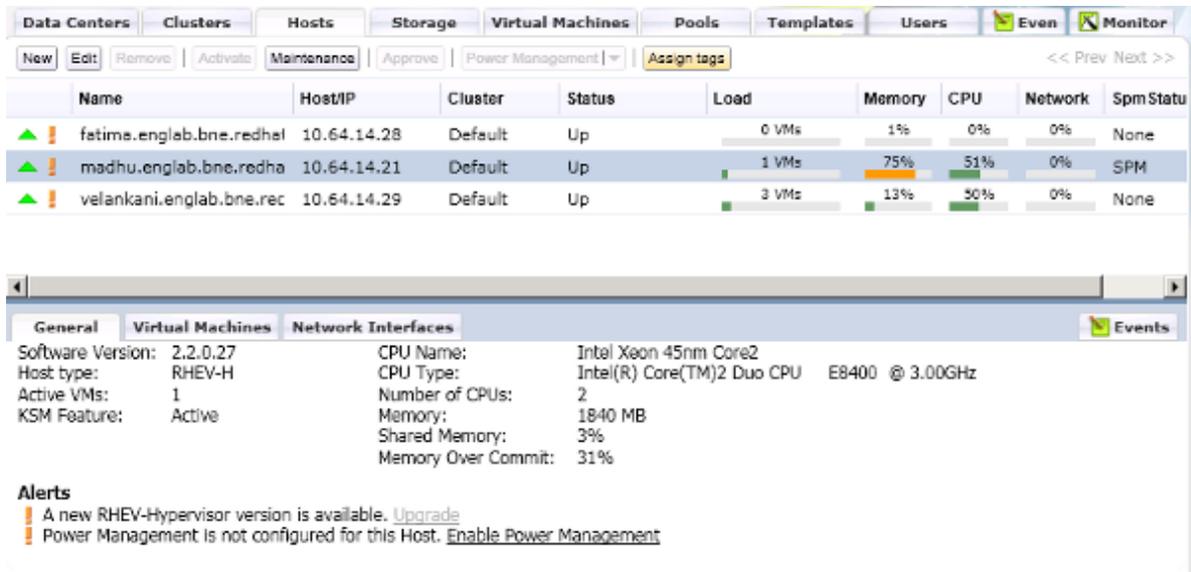
Figure 7.1. General Pane with Upgrade Button

An Alert Message indicates that a new version of the RHEV Hypervisor is available. The **Upgrade** link is disabled if the host is has a status of **Up**. A tooltip directs you to switch to maintenance mode to enable upgrade.
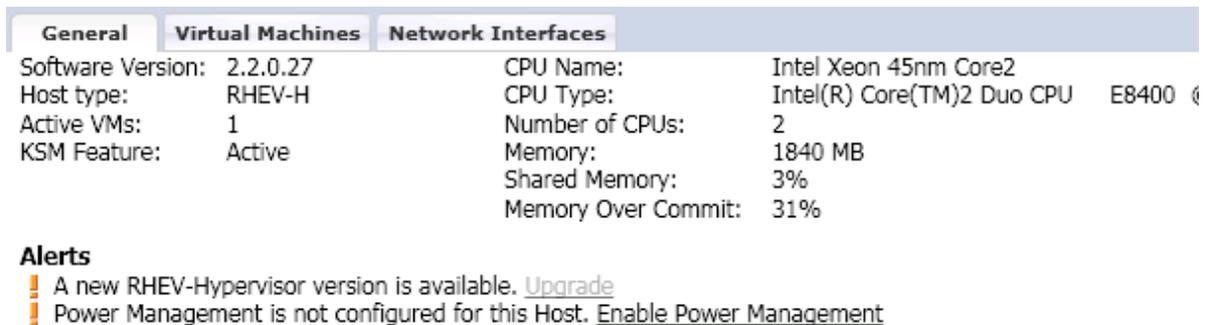
Figure 7.2. Upgrade Button

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 selectable link to re-install or upgrade the host.

4. Click the link. The **Install Host** dialog displays.
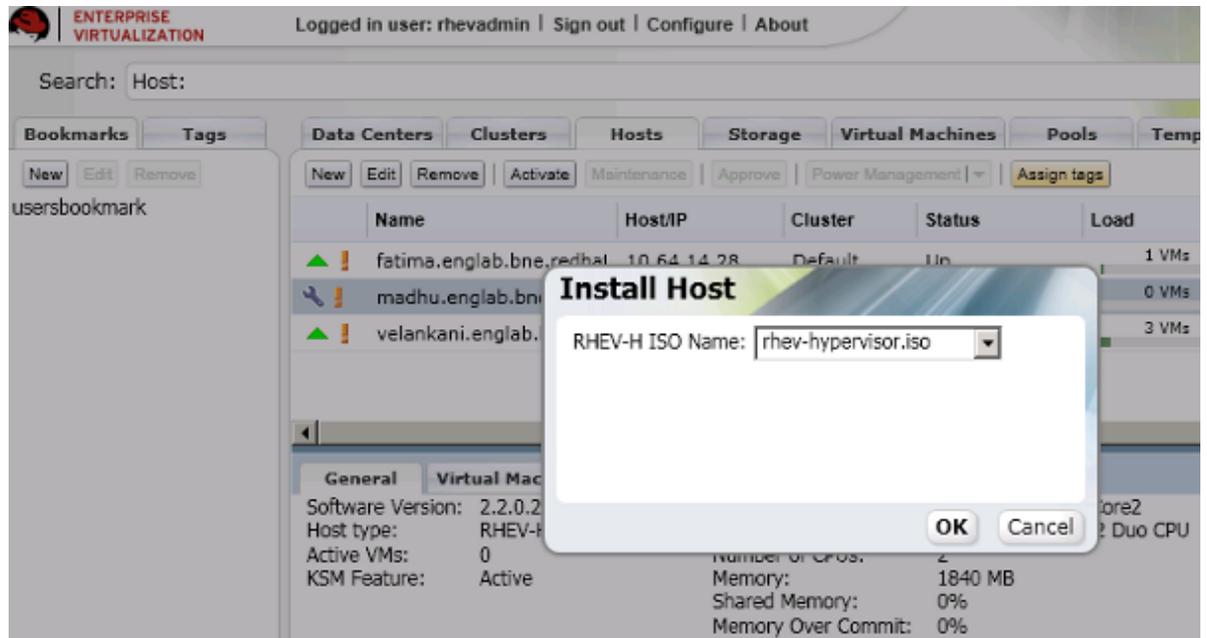
Figure 7.3. Install Host Dialog

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 normal, and may take some time.

7. Once successfully upgraded, the host displays a status of **Up**. Any virtual machines that were migrated off the host, may be migrated back onto the upgraded host.

7.2. Upgrading a RHEV Hypervisor with local media

To upgrade an existing hypervisor from a local CD-ROM or USB device, insert the installation media and ensure the system boots from it. Boot the kernel with the upgrade parameter. This will automatically upgrade and reboot the system, rather than taking you to the interactive configuration menu. For more information, refer to *upgrade*.

```
ISOLINUX 3.11 2005-09-02  Copyright (C) 1994-2005 H. Peter Anvin
This is the Red Hat Enterprise Virtualization Hypervisor 5.7-2.2 (4.2)
- to boot press the <ENTER> key
or type linux and press the <ENTER> key.
- Extra kernel command line parameters can be passed following the
linux keyword.
- Boot parameters summary:
BOOTIF=link|eth*|<MAC>(pxelinux format)
storage_init=[usb|scsi|cciss[:serial#]]|/dev/...
storage_vol=:SWAP_MB::CONFIG_MB:LOGGING_MB:DATA_MB
mem_overcommit=<overcommit_ratio>
upgrade
ip=<client-ip> netmask=<netmask> gateway=<gw-ip>
vlan=<VLAN-id>
dns=server[:server]
ntp=server[:server]
hostname=fqdn
management_server=RHEVM-host[:port]
netconsole_server=RHEVM-host[:port]
boot: linux upgrade
```

Figure 7.4. Booting with the upgrade parameter

7.3. Re-installing RHEV Hypervisors with the RHEV Manager

Re-installing RHEV Hypervisors is the same procedure as upgrading, refer to *Procedure 7.1, "Upgrading a RHEV Hypervisor"* for details on upgrading.

To reset settings, remove the previous Hypervisor first with the **Uninstall an existing Hypervisor** menu item and then re-install the Hypervisor.

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 RHEV products by email. Subscribe by completing this form: *http://www.redhat.com/mailman/listinfo/rhev-watch-list/*.

RHEV 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 RHEV. The list below covers ports and their usage by RHEV:

- ICMP requests must be accepted. ICMP packets are used for network testing by the RHEV Manager.

- Port 22 should be open for SSH access and the initial installation.

- Ports 80 or 443 (depending on the security settings on the RHEV 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 A. Uninstalling Red Hat Enterprise Virtualization Hypervisors from local storage devices

This appendix covers uninstalling the Red Hat Enterprise Virtualization (RHEV) Hypervisors from local storage devices.

Procedure A.1. Removing RHEV Hypervisors

1. Boot the RHEV Hypervisor. Refer to *Chapter 4, Preparing Red Hat Enterprise Virtualization Hypervisor installation media* for details on booting Red Hat Enterprise Virtualization Hypervisors.

    The `firstboot` kernel parameter is required for uninstalling hypervisors.

    ```
    : linux firstboot
    ```

2. Select **Support Menu** from the **Hypervisor Configuration Menu**. From the **Support Menu** select **Uninstall an existing Hypervisor**.

3. The following message will appear:

    ```
    !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
    !!                                                                         !!
    !! WARNING !! WARNING !! WARNING !! WARNING !! WARNING !! WARNING !! WARNING !!
    !!                                                                         !!
    !!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!

       If you proceed, you will DESTROY any existing data related to the RHEV
       installation on this machine.

    Do you wish to continue and uninstall this Hypervisor (Y/N)?
    ```

    Press **Y** to confirm and remove the Hypervisor.

The RHEV Hypervisor is now removed.

Appendix B. Revision History

Revision 4-0 **Thursday June 30 2011** **Stephen Gordon** *sgordon@redhat.com*
>
> Updated instructions for boot media creation. Tools are now included in the Hypervisor ISO itself, BZ#690254.

Revision 3-0 **Tuesday July 6 2010** **David Jorm** *djorm@redhat.com*
>
> Documented required parameters for automated installation, BZ#611440

Revision 2-0 **Thursday June 24 2010** **David Jorm** *djorm@redhat.com*
>
> Fixed rhn_url parameter example, BZ#607274
> Updated boot menu text to show release build banner, BZ#594699

Revision 1-0 **Thursday June 10 2010** **David Jorm** *djorm@redhat.com*
>
> RHEV 2.2
> General Availability