



Red Hat Enterprise Linux 8

Performing a standard RHEL installation

Installation documentation for the release of Red Hat Enterprise Linux

Red Hat Enterprise Linux 8 Performing a standard RHEL installation

Installation documentation for the release of Red Hat Enterprise Linux

Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

Abstract

This document is for users who want to install Red Hat Enterprise Linux using the graphical user interface.

Table of Contents

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION	6
CHAPTER 1. INTRODUCTION	7
1.1. SUPPORTED ARCHITECTURES	7
CHAPTER 2. INSTALLATION METHODS	8
Additional resources	8
2.1. PERFORMING A QUICK INSTALL ON AMD, INTEL 64-BIT, AND 64-BIT ARM	8
Prerequisites	8
Procedure	8
Additional resources	9
PART I. PERFORMING A CUSTOM INSTALL ON AMD, INTEL 64-BIT, AND 64-BIT ARM	10
CHAPTER 3. INSTALLATION WORKFLOW	11
Additional resources	11
CHAPTER 4. PREPARING FOR YOUR INSTALLATION	12
4.1. RECOMMENDED STEPS	12
4.2. CHECK SYSTEM REQUIREMENTS	12
Additional resources	12
4.3. CHOOSE AN INSTALLATION BOOT METHOD	12
Additional Resources	13
4.4. SELECT THE REQUIRED INSTALLATION IMAGE	13
Additional Resources	14
4.5. DOWNLOAD THE INSTALLATION ISO IMAGE	14
4.5.1. Downloading an ISO image from the Customer Portal	14
4.5.2. Downloading an ISO image using curl	15
Prerequisites	15
Procedure	15
4.6. CREATE INSTALLATION MEDIA	16
4.6.1. Creating a bootable CD or DVD	16
4.6.2. Creating a bootable USB device on Linux	16
4.6.3. Creating a bootable USB device on Windows	17
Prerequisites	17
Procedure	18
4.6.4. Creating a bootable USB device on Mac OS X	18
Prerequisites	18
Procedure	19
4.7. PREPARE AN INSTALLATION SOURCE	20
4.7.1. Types of installation source	20
4.7.2. Specify an installation source	21
4.7.3. Creating an installation source on an NFS server	21
Prerequisites	21
Procedure	22
4.7.4. Creating an installation source using HTTP or HTTPS	23
Prerequisites	23
Procedure	23
Additional resources	24
4.7.5. Creating an installation source using FTP	24
Prerequisites	24
Procedure	24

CHAPTER 5. BOOTING THE INSTALLATION	27
5.1. BOOTING THE INSTALLATION FROM A USB, CD, OR DVD	27
5.2. BOOTING THE INSTALLATION FROM A NETWORK USING PXE	27
Prerequisites	27
Procedure	28
5.3. BOOT OPTIONS REFERENCE	28
CHAPTER 6. INSTALLING USING THE GRAPHICAL USER INTERFACE	29
6.1. GRAPHICAL INSTALLATION WORKFLOW	29
6.2. CONFIGURING LANGUAGE AND LOCATION SETTINGS	29
6.3. INSTALLATION SUMMARY	30
6.4. LOCALIZATION SETTINGS	32
6.4.1. Configuring keyboard, language, and time and date settings	32
6.5. SOFTWARE SETTINGS	34
6.5.1. Configuring installation source	34
6.5.2. Configuring software selection	36
6.6. SYSTEM SETTINGS	37
6.6.1. Configuring installation destination	37
Procedure	38
6.6.1.1. Configuring boot loader	41
6.6.2. Configuring Kdump	42
Procedure	42
6.6.3. Network and host name	43
6.6.3.1. Configuring network and host name	43
6.6.3.2. Adding a virtual network interface	43
6.6.3.3. Editing network interface configuration	44
6.6.3.4. Enabling or Disabling the Interface Connection	45
6.6.3.5. Setting up Static IPv4 or IPv6 Settings	45
6.6.3.6. Configuring Routes	46
6.6.4. Security policy	46
6.6.4.1. About security policy	46
6.6.4.2. Configuring a security policy	47
6.6.4.3. Related information	47
6.6.5. System Purpose	48
Additional resources	49
6.6.5.1. Configuring system purpose using the graphical user interface	49
6.6.5.2. Configuring System Purpose using Kickstart	49
Additional resources	50
6.7. STORAGE DEVICES	50
6.7.1. Storage device selection	51
6.7.2. Filtering storage devices	51
6.7.3. Advanced storage options	52
6.7.3.1. Discovering and starting an iSCSI session	53
6.7.3.2. Configuring FCoE parameters	54
6.7.3.3. Configuring DASD storage devices	55
6.7.3.4. Configuring FCP devices	55
6.7.4. Installing to an NVDIMM device	56
6.7.4.1. Criteria for using an NVDIMM Device as an installation target	56
6.7.4.2. Configuring an NVDIMM device using Anaconda	57
6.7.4.3. Configuring an NVDIMM device using kickstart	58
6.8. MANUAL PARTITIONING	59
6.8.1. Starting manual partitioning	59
6.8.2. Adding a mount point file system	61

6.8.3. Configuring a mount point file system	61
6.8.4. Customizing a partition or volume	62
6.8.5. Creating software RAID	64
6.8.6. Creating an LVM logical volume	65
6.8.7. Configuring an LVM logical volume	65
6.9. STARTING THE INSTALLATION PROGRAM	66
6.9.1. Beginning installation	66
6.9.2. Configuring a root password	67
6.9.3. Creating a user account	68
6.9.3.1. Configuring Advanced User Settings	69
6.9.4. Installation complete	70
6.9.5. UEFI Secure Boot for RHEL 8	70
6.9.6. Adding a custom private key for UEFI Secure Boot	70
CHAPTER 7. POST-INSTALLATION TASKS	72
7.1. COMPLETING INITIAL SETUP	72
Prerequisite	72
Procedure	72
Additional resources	73
7.2. REGISTERING YOUR SYSTEM USING THE COMMAND LINE	73
7.3. REGISTERING YOUR SYSTEM USING THE SUBSCRIPTION MANAGER USER INTERFACE	74
Prerequisites	74
Procedure	74
7.4. REGISTRATION ASSISTANT	75
7.5. SECURING YOUR SYSTEM	75
Prerequisites	75
Procedure	75
Additional resources	76
APPENDIX A. SYSTEM REQUIREMENTS REFERENCE	77
A.1. CHECK HARDWARE COMPATIBILITY	77
A.2. REVIEW SUPPORTED INSTALLATION TARGETS	77
A.3. RECORD SYSTEM SPECIFICATIONS	77
A.4. CHECK DISK AND MEMORY REQUIREMENTS	78
A.5. REVIEW RAID REQUIREMENTS	79
APPENDIX B. PARTITIONING REFERENCE	81
B.1. SUPPORTED DEVICE TYPES	81
B.2. SUPPORTED FILE SYSTEMS	81
B.3. SUPPORTED RAID TYPES	82
B.4. RECOMMENDED PARTITIONING SCHEME	82
B.5. ADVICE ON PARTITIONS	85
APPENDIX C. TROUBLESHOOTING	87
C.1. CONSOLES AND LOGGING DURING INSTALLATION	87
C.2. SAVING SCREENSHOTS	88
C.3. RESUMING AN INTERRUPTED DOWNLOAD ATTEMPT	88
PART II. INSTALLING RED HAT ENTERPRISE LINUX ON IBM POWER	89
CHAPTER 8. PREPARATION FOR IBM POWER SYSTEMS SERVERS	90
8.1. PREREQUISITES	90
8.2. INSTALLATION OR UPGRADING	90
8.2.1. Planning for Installation on IBM Power Systems	90
Additional Resources	90

8.3. PREPARATION FOR IBM POWER SYSTEMS SERVERS	90
Additional Resources	90
8.4. IBM INSTALLATION TOOLS	91
8.4.1. IBM Installation Toolkit	91
8.5. POWERLINUX PRODUCTIVITY TOOLS	91
Additional resources	91
8.6. SUPPORTED INSTALLATION TARGETS	91
Additional resources	92
8.7. SYSTEM SPECIFICATIONS LIST	92
8.8. DISK SPACE AND MEMORY REQUIREMENTS	93
Additional Resources	94
8.9. RAID AND OTHER DISK DEVICES	94
8.9.1. RAID and Other Disk Devices	94
8.9.1.1. Hardware RAID	94
8.9.1.2. Software RAID	94
8.9.1.3. USB Disks	94
Additional Resources	95
8.10. CHOOSE AN INSTALLATION BOOT METHOD	95
Additional Resources	95
8.11. AUTOMATING THE INSTALLATION WITH KICKSTART	95
Additional Resources	96
8.12. RELATED INFORMATION	96
PART III. INSTALLING RED HAT ENTERPRISE LINUX ON IBM Z	97
CHAPTER 9. PREPARING FOR INSTALLATION ON IBM Z	98
9.1. PREREQUISITES	98
9.2. OVERVIEW OF THE IBM Z INSTALLATION PROCESS	98
Additional resources	98
9.3. PLANNING FOR INSTALLATION ON IBM Z	98
9.3.1. Pre-installation	98
Additional Resources	99
9.4. INSTALLING UNDER Z/VM	100
9.5. USING PARAMETER AND CONFIGURATION FILES ON IBM Z	101
9.6. REQUIRED CONFIGURATION FILE PARAMETERS ON IBM Z	101
9.7. IBM ZVM CONFIGURATION FILE	102
9.8. INSTALLATION NETWORK PARAMETERS ON IBM Z	102
9.9. PARAMETERS FOR KICKSTART INSTALLATIONS ON IBM Z	105
9.10. MISCELLANEOUS PARAMETERS ON IBM Z	106
9.11. SAMPLE PARAMETER FILE AND CMS CONFIGURATION FILE ON IBM Z	107
CHAPTER 10. CONFIGURING A LINUX INSTANCE ON IBM Z	108
10.1. PREREQUISITES	108
10.2. DASDS THAT ARE PART OF THE ROOT FILE SYSTEM	108
10.3. FCP LUNS THAT ARE PART OF THE ROOT FILE SYSTEM	110
10.4. ADDING A QETH DEVICE	112
10.5. DYNAMICALLY ADDING A QETH DEVICE	112
10.6. PERSISTENTLY ADDING A QETH DEVICE	114
10.7. CONFIGURING AN IBM Z NETWORK DEVICE FOR NETWORK ROOT FILE SYSTEM	117

PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so:

- For simple comments on specific passages, make sure you are viewing the documentation in the Multi-page HTML format. Highlight the part of text that you want to comment on. Then, click the **Add Feedback** pop-up that appears below the highlighted text, and follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
 1. Go to the [Bugzilla](#) website.
 2. As the Component, use **Documentation**.
 3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
 4. Click **Submit Bug**.

CHAPTER 1. INTRODUCTION

Red Hat Enterprise Linux 8 delivers a stable, secure, consistent foundation across hybrid cloud deployments with the tools needed to deliver workloads faster with less effort. It can be deployed as a guest on supported hypervisors and Cloud provider environments as well as deployed on physical infrastructure, so your applications can take advantage of innovations in the leading hardware architecture platforms.

1.1. SUPPORTED ARCHITECTURES

Red Hat Enterprise Linux supports the following architectures:

- AMD, Intel 64-bit, and 64-bit ARM
- IBM Power Systems
- IBM Z

CHAPTER 2. INSTALLATION METHODS

There are three methods to install Red Hat Enterprise Linux:

Quick

Install Red Hat Enterprise Linux on AMD, Intel 64-bit, and 64-bit ARM architectures using the graphical installation. The quick installation assumes that you accept the default settings provided by the installation program.

Custom

Install Red Hat Enterprise Linux on all architectures using the graphical installation. The custom installation allows you to configure the graphical installation settings for your specific requirements.

Automated

Install Red Hat Enterprise Linux on all architectures using Kickstart. An automated installation allows you to perform unattended operating system installation tasks.

Additional resources

- To perform a custom install on AMD, Intel 64-bit, and 64-bit ARM architectures using the graphical user interface, see [link to be added](#)
- To perform a custom install on IBM Power using the graphical user interface, see [link to be added](#)
- To perform a custom install on IBM Z using the graphical user interface, see [link to be added](#)
- To perform an automated install, see the *Performing an advanced RHEL installation* guide.

2.1. PERFORMING A QUICK INSTALL ON AMD, INTEL 64-BIT, AND 64-BIT ARM

This procedure describes how to perform a quick installation on AMD, Intel 64-bit, and 64-bit ARM architectures using the graphical user interface. This procedure assumes that you are familiar with Red Hat Enterprise Linux and your environment, and can accept the default settings provided by the installation program.

Prerequisites

- You have downloaded the required ISO image file and created a bootable device.
- You have booted the installation program and the boot menu is displayed.

Procedure

1. From the boot menu, select **Install Red Hat Enterprise Linux 8.0**.
2. Press the **Enter** key on your keyboard.
3. From the **Welcome to Red Hat Enterprise Linux 8.0** window, select your language and location.
4. Click **Continue** to proceed to the **Installation Summary** window.

**NOTE**

The **Installation Summary** window is the central hub from which installation settings are configured. The default settings assigned by the installation program are displayed under each category.

5. From the **Installation Summary** window, select **Installation Destination**.
 - a. From the **Local Standard Disks** pane, select the target disk.
 - b. Click **Done** to accept the selection and the default setting of automatic partitioning, and return to the **Installation Summary** window.
6. From the **Installation Summary** window, select **Network & Host Name**.
 - a. Toggle the **Ethernet** switch to **ON** to enable network configuration.
 - i. Optional: Select a network device and click **Configure** to update the network interface configuration.
 - b. Click **Done** to accept the changes and return to the **Installation Summary** window.
7. From the **Installation Summary** window, select **Security Policy**.
 - a. Select the required profile and click **Select profile**.
 - b. Click **Done** to accept the changes and return to the **Installation Summary** window.
8. From the **Installation Summary** window, select **System Purpose**.
 - a. Select the required role, service level agreement, and usage.
 - b. Click **Done** to accept the changes and return to the **Installation Summary** window.
9. From the **Installation Summary** window, click **Begin Installation** to start the installation.
10. From the **Configuration** window, configure a root password and create a user account.
11. Once the installation completes, click **Reboot** to restart the system.
12. From the **Initial Setup** window, accept the licensing agreement and register your system.

Additional resources

- To download an ISO image file, see **link to be added**.
- To boot the installation program, see **link to be added**.
- To customize the graphical installation settings, see **link to be added**.
- To register your system, see **link to be added**.

PART I. PERFORMING A CUSTOM INSTALL ON AMD, INTEL 64-BIT, AND 64-BIT ARM

This section describes how to perform a custom installation of Red Hat Enterprise Linux on AMD, Intel 64-bit, and 64-bit ARM architectures using the graphical user interface.

CHAPTER 3. INSTALLATION WORKFLOW

This installation workflow is for users performing a customized installation on AMD, Intel 64-bit, and 64-bit ARM architectures using the graphical user interface.

1. Prepare for your installation by checking your system and hardware requirements, downloading an installation image file, and creating installation media.
2. Boot the installation program and install Red Hat Enterprise Linux using the graphical user interface.
3. Complete post-installation tasks such as initial setup and system registration.

Additional resources

- To prepare for your installation, see [Chapter 4, *Preparing for your installation*](#).
- To boot the installation program, see [Chapter 5, *Booting the installation*](#).
- To install using the graphical user interface, see [Chapter 6, *Installing using the Graphical User Interface*](#)
- To complete post-installation tasks, see [Chapter 7, *Post-installation tasks*](#).

CHAPTER 4. PREPARING FOR YOUR INSTALLATION

If you are new to Red Hat Enterprise Linux, it is important to prepare for your installation by reviewing system requirements, downloading the required installation image, and creating installation media.

4.1. RECOMMENDED STEPS

Preparing for your installation consists of several steps.



NOTE

- If you are new to Red Hat Enterprise Linux, complete steps 1 to 5.
- If you are familiar with Red Hat Enterprise Linux, complete steps 3 to 5.

Steps

1. Check system requirements.
2. Choose an installation boot method.
3. Select and download the installation image.
4. Create the installation media.
5. Prepare the installation source*

*Only required for the boot ISO (minimal install) image.

4.2. CHECK SYSTEM REQUIREMENTS


If this is a first-time install of Red Hat Enterprise Linux it is recommended that you review the guidelines provided for system, hardware, security, memory, and RAID before installing. See [Appendix A, System requirements reference](#) for more information.

Additional resources

- For more information on securing Red Hat Enterprise Linux, see the *Configuring and managing security* guide.

4.3. CHOOSE AN INSTALLATION BOOT METHOD

There are several methods to boot the Red Hat Enterprise Linux installation program. The method you choose depends upon your installation media.



NOTE

Installation media must remain mounted throughout the installation process, including during the execution of the **%post** section of a kickstart file.

Full installation DVD or USB drive

A full installation DVD or USB drive is created using the Binary DVD ISO image. It can be used as both a boot device and as an installation source for installing software packages.

Minimal installation CD, DVD or USB flash drive

A minimal installation CD, DVD, or USB flash drive is created using the boot ISO image, which only contains the minimum files necessary to boot the system and start the installation program. This boot option requires an installation source that contains the required software packages.

PXE Server

A *preboot execution environment* (PXE) server allows the installation program to boot over the network. After you boot the system, you complete the installation from a different installation source, such as a local hard drive or a location on a network.

Additional Resources

- For instructions on how to create a full installation DVD or USB drive, see [Section 4.6, “Create installation media”](#) for more information.
- For instructions on how to create a boot CDs, DVDs and USB flash drive, see [Section 4.7, “Prepare an installation source”](#) for more information.
- For more information on PXE servers, see <XX TO BE ADDED>

4.4. SELECT THE REQUIRED INSTALLATION IMAGE

Two types of Red Hat Enterprise Linux 8 installation images are available.

Binary DVD ISO image file

A full installation program that contains the BaseOS and AppStream repositories and allows you to complete the installation without additional repositories.



IMPORTANT

- It is **recommended** that the Binary DVD ISO image file is used to install Red Hat Enterprise Linux 8.
- A Binary DVD for IBM Z can be used to boot the installation program using a SCSI DVD drive, or as an installation source.

Boot ISO image file

The Boot ISO image is a minimal installation that requires access to the BaseOS and AppStream repositories to install software packages. The repositories are part of the Binary DVD ISO image that is available for download from <https://access.redhat.com/home>. Download and unpack the Binary DVD ISO image to access the repositories.

The following table lists the images available for the supported architectures.

Table 4.1. Boot and Installation Images

Architecture	Installation DVD	Boot DVD
AMD64 and Intel 64	x86_64 Binary DVD ISO image file	x86_64 Boot ISO image file
ARM 64	AArch64 Binary DVD ISO image file	AArch64 Boot ISO image file

Architecture	Installation DVD	Boot DVD
IBM POWER	ppc64le Binary DVD ISO image file	ppc64le Boot ISO image file
IBM Z	s390x Binary DVD ISO image file	s390x Boot ISO image file

Additional Resources

- For instructions on how to access the Binary DVD ISO image repositories, see [Section 4.7, “Prepare an installation source”](#) for more information.

4.5. DOWNLOAD THE INSTALLATION ISO IMAGE

This section provides instructions on how to download a Red Hat Enterprise Linux installation image from the Red Hat Customer Portal or by using the **curl** command.

4.5.1. Downloading an ISO image from the Customer Portal

This procedure describes how to download a Red Hat Enterprise Linux 8 ISO image from the Red Hat Customer Portal.



NOTE

- The Binary DVD ISO image is the recommended method for installing Red Hat Enterprise Linux 8 as it contains all repositories and software packages, and does not require any additional configuration.
- If you download the Boot ISO image file, you must configure an installation source to obtain the repositories and software packages. See [Section 4.7, “Prepare an installation source”](#) for more information.

Prerequisites

- You have an active Red Hat subscription.
- You are logged in to the Red Hat Customer Portal at <https://access.redhat.com/home>.

Procedure

- Click **DOWNLOADS**.
- Select the **By Category** tab.
- Click the **Red Hat Enterprise Linux** link.
The **Download Red Hat Enterprise Linux** web page opens.
- From the **Product Variant** drop-down menu, select the required variant.

**NOTE**

If you are unsure of the variant for your requirements, see <http://www.redhat.com/en/technologies/linux-platforms/enterprise-linux>.

5. From the **Version** drop-down menu, select **8.0**.
6. From the **Architecture** drop-down menu, select the required architecture.
A **Binary DVD** and a **Boot ISO** image file are displayed under the **Product Software** tab. Additional images may be available, for example, preconfigured virtual machine images, but they are beyond the scope of this document.
7. Click **Download Now** beside the required image.

4.5.2. Downloading an ISO image using curl

Use the **curl** command to download installation images directly from a specific URL.

Prerequisites

1. Verify the curl package is installed:

- If your distribution uses the **yum** package manager:

```
# yum install curl
```

- If your distribution uses the **dnf** package manager:

```
# dnf install curl
```

- If your distribution uses the **apt** package manager:

```
# apt update
# apt install curl
```

- On other systems, download curl from the curl website. If your Linux distribution does not use yum, dnf, or apt, or if you do not use Linux, download the most appropriate software package from the [curl web site](http://curl.haxx.se).

Procedure

1. On the command line, enter a suitable directory, and type the following command to download the file:

```
$ curl --output directory-path/filename.iso 'copied_link_location'
```

2. Replace *directory-path* with a path where you want to save the file; replace *filename.iso* with the ISO image name as displayed in the Customer Portal; replace *copied_link_location* with the link that you have copied from the Customer Portal.

Example:

```
$ curl --output Downloads/rhel-server-8.0-x86_64-dvd.iso  
'https://access.redhat.com/downloads/content/69/ver=/rhel---  
8/8.0/x86_64/product-software
```

4.6. CREATE INSTALLATION MEDIA

This section describes how to use an ISO image file downloaded in [Section 4.5, “Download the installation ISO image”](#) to create bootable physical media, such as a DVD or a USB flash drive. These steps are for installing Red Hat Enterprise Linux on an AMD, Intel 64-bit, or 64-bit ARM system.



NOTE

By default, the `inst.stage2=boot` option is used on the installation media and is set to a specific label, for example, `inst.stage2=hd:LABEL=RHEL8\x20Server.x86_64`. If you modify the default label of the file system containing the runtime image, or if you use a customized procedure to boot the installation system, you must verify the label is set to the correct value.

4.6.1. Creating a bootable CD or DVD

The steps to create a bootable CD or DVD from an ISO image file vary, depending on the operating system and disc burning software. Consult your burning software’s documentation for the steps needed to burn a CD or DVD from an ISO image file.



WARNING

The Binary DVD ISO image may be larger than 4.7 GB and as a result, it may not fit on a single-layer DVD. A dual-layer DVD or USB key is recommended.

4.6.2. Creating a bootable USB device on Linux

This procedure describes how to create a bootable USB device on a Linux system.

Prerequisites

- You have downloaded an installation ISO image as described in [Section 4.5, “Download the installation ISO image”](#).
- You have a USB flash drive that is large enough to hold the ISO image.



NOTE

This procedure is destructive and data on the USB flash drive is destroyed without warning.

Procedure

1. Connect the USB flash drive to the system.

2. Open a terminal window and run the **dmesg** command. A log detailing all recent events is displayed. At the bottom of this log, messages resulting from the attached USB flash drive are displayed. Record the name of the connected device.

3. Log in as root:

```
$ su -
```

4. Enter your root password when prompted.

5. Find the device node assigned to the drive. In this example, the drive is given the name **sdd**.

```
$ dmesg|tail
[288954.686557] usb 2-1.8: New USB device strings: Mfr=0, Product=1,
SerialNumber=2
[288954.686559] usb 2-1.8: Product: USB Storage
[288954.686562] usb 2-1.8: SerialNumber: 000000009225
[288954.712590] usb-storage 2-1.8:1.0: USB Mass Storage device
detected
[288954.712687] scsi host6: usb-storage 2-1.8:1.0
[288954.712809] usbcore: registered new interface driver usb-storage
[288954.716682] usbcore: registered new interface driver uas
[288955.717140] scsi 6:0:0:0: Direct-Access      Generic  STORAGE
DEVICE    9228 PQ: 0 ANSI: 0
[288955.717745] sd 6:0:0:0: Attached scsi generic sg4 type 0
[288961.876382] sd 6:0:0:0: sdd Attached SCSI removable disk
```

6. Use the **dd** command to write the ISO image directly to the USB device.

```
# dd if=/image_directory/image.iso of=/dev/device
```

7. Replace */image_directory/image.iso* with the full path to the ISO image file downloaded
8. Replace *device* with the device name as reported by the **dmesg** command; in this example it would be **sdd**.
9. Verify the output as the device name, for example, */dev/sdd*, not as a name of a partition on the device, for example, */dev/sdd1*. If the ISO image is located in */home/testuser/Downloads/rhel-server-8-x86_64-boot.iso* and the detected device name is **sdd**, the command is:

```
$ dd if=/home/testuser/Downloads/rhel-server-8-x86_64-boot.iso
of=/dev/sdd
```

10. Wait for the **dd** command to finish writing the image to the device. The data transfer is finished when the **#** prompt appears. Once the prompt is displayed, log out of the root account and unplug the USB drive. The USB drive is now ready to be used as a boot device.

4.6.3. Creating a bootable USB device on Windows

This procedure describes how to create a bootable USB device on a Windows system. The procedure varies depending on the tool. Red Hat recommends using Fedora Media Writer, available for download at <https://github.com/MartinBriza/MediaWriter/releases>.

Prerequisites

- You have downloaded an installation ISO image as described in [Section 4.5, “Download the installation ISO image”](#).
- You have a USB flash drive that is large enough to hold the ISO image.

**NOTE**

This procedure is destructive and data on the USB flash drive is destroyed without warning.

Procedure

1. Download and install Fedora Media Writer from <https://github.com/MartinBriza/MediaWriter/releases>

**NOTE**

To install Fedora Media Writer on Red Hat Enterprise Linux use the pre-built Flatpak package. The package can be obtained from the official Flatpak repository Flathub.org at <https://flathub.org/apps/details/org.fedoraproject.MediaWriter>

2. Connect the USB flash drive to the system.
3. Open Fedora Media Writer.
4. From the main window, click **Custom Image** and select the previously downloaded Red Hat Enterprise Linux ISO image.
5. From **Write Custom Image** window, select the drive you want to use.
6. Click **Write to disk**. The boot media creation process starts. Do not unplug the drive until the operation completes. The operation may take several minutes, depending on the size of the ISO image, and the write speed of the USB drive.
7. When the operation completes, unmount the USB drive. The USB drive is now ready to be used as a boot device.

4.6.4. Creating a bootable USB device on Mac OS X

This procedure describes how to create a bootable USB device on a Mac OS X system.

Prerequisites

- You have downloaded an installation ISO image as described in [Section 4.5, “Download the installation ISO image”](#).
- You have a USB flash drive that is large enough to hold the ISO image.

**NOTE**

This procedure is destructive and data on the USB flash drive is destroyed without warning.

Procedure

1. Connect the USB flash drive to the system.
2. Identify the device path with the **diskutil list** command. The device path has the format of */dev/disknumber*, where number is the number of the disk. The disks are numbered starting at zero (0). Typically, Disk 0 is the OS X recovery disk, and Disk 1 is the main OS X installation. In the following example, it is **disk2**:

```
$ diskutil list
/dev/disk0
#:                                TYPE NAME                                SIZE
IDENTIFIER
0:      GUID_partition_scheme                                *500.3 GB
disk0
1:                                EFI EFI                                209.7 MB
disk0s1
2:      Apple_CoreStorage                                400.0 GB
disk0s2
3:      Apple_Boot Recovery HD                                650.0 MB
disk0s3
4:      Apple_CoreStorage                                98.8 GB
disk0s4
5:      Apple_Boot Recovery HD                                650.0 MB
disk0s5
/dev/disk1
#:                                TYPE NAME                                SIZE
IDENTIFIER
0:      Apple_HFS YosemiteHD                                *399.6 GB
disk1
Logical Volume on disk0s1
8A142795-8036-48DF-9FC5-84506DFBB7B2
Unlocked Encrypted
/dev/disk2
#:                                TYPE NAME                                SIZE
IDENTIFIER
0:      FDisk_partition_scheme                                *8.0 GB
disk2
1:      Windows_NTFS SanDisk USB                                8.0 GB
disk2s1
```

3. To identify your USB flash drive, compare the NAME, TYPE and SIZE columns to your flash drive. For example, the NAME should be the title of the flash drive icon in the **Finder** tool. You can also compare these values to those in the flash drive's information panel.
4. Use the **diskutil unmountDisk** command to unmount the flash drive's filesystem volumes:

```
$ diskutil unmountDisk /dev/disknumber
Unmount of all volumes on disknumber was successful
```

When the command completes, the icon for the flash drive disappears from your desktop. If it does not, you may have selected the wrong disk. Attempting to unmount the system disk accidentally, returns a **failed to unmount** error.

5. Log in as root:

```
$ su -
```

6. Enter your root password when prompted.
7. Use the **dd** command as a parameter of the **sudo** command to write the ISO image to the flash drive:

```
$ sudo dd if=/path/to/image.iso of=/dev/rdisknumber bs=1m>
```



NOTE

Mac OS X provides both a block (`/dev/disk*`) and character device (`/dev/rdisk*`) file for each storage device. Writing an image to the `/dev/rdisknumber` character device is faster than to the `/dev/disknumber` block device.


8. To write the `/Users/user_name/Downloads/rhel-server-8-x86_64-boot.iso` file to the `/dev/rdisk2` device:

```
$ sudo dd if=/Users/user_name/Downloads/rhel-server-8-x86_64-  
boot.iso of=/dev/rdisk2
```

9. Wait for the **dd** command to finish writing the image to the device. The data transfer is finished when the **#** prompt appears. Once the prompt is displayed, log out of the root account and unplug the USB drive. The USB drive is now ready to be used as a boot device.

4.7. PREPARE AN INSTALLATION SOURCE

This section describes how to create an installation source for the Boot ISO image using the Binary DVD ISO image that contains the required repositories and software packages. The Boot ISO image file does not include any repositories or software packages; it only contains the installation program and the tools required to boot the system and start the installation.



IMPORTANT

Creating an installation source is only required for the Boot ISO image. Red Hat recommends the Binary DVD ISO image as the preferred method to install Red Hat Enterprise Linux.

4.7.1. Types of installation source

The installation source for a minimal boot image can be a:

- **DVD:** Burn the Binary DVD ISO image to a DVD and configure the installation program to install the software packages from the DVD.
- **Hard drive or USB drive:** Copy the Binary DVD ISO image and configure the installation program to install the software packages from the drive. If you use a USB drive, verify that it is connected to the system before the installation begins. The installation program is not able to detect media after the installation begins.

**NOTE**

A limitation of using a hard drive as the installation source is that the Binary DVD ISO image on the hard drive must be on a partition with a file system that the installation program can mount. The file systems are xfs, ext2, ext3, ext4, and vfat (FAT32). On Microsoft Windows systems, the default file system used when formatting hard drives is NTFS, and the exFAT file system is also available. However, neither of these file systems can be mounted during the installation. If you are creating a hard drive or a USB drive as an installation source on Microsoft Windows, verify that you formatted the drive as FAT32. The FAT32 file system does not support files larger than 4 GiB.

- **Network location:** Copy the Binary DVD ISO image or the installation tree (extracted contents of the Binary DVD ISO image) to a network location and perform the installation over the network using the following protocols:
 - **NFS:** The Binary DVD ISO image is in a Network File System (NFS) share.
 - **HTTPS, HTTP or FTP:** The installation tree is on a network location that is accessible over HTTP, HTTPS, or FTP.

4.7.2. Specify an installation source

You can specify an installation source using any of the following methods:

- **Graphical installation:** Select the installation source in the **Installation Source** window of the graphical install. See [Section 6.5.1, “Configuring installation source”](#) for more information.
- **Boot option:** Configure a custom boot option to specify the installation source. See [Section 5.3, “Boot options reference”](#) for more information.
- **Kickstart file:** Use the install command in a Kickstart file to specify the installation source. See the *Installing RHEL as an experienced user* guide for more information.

4.7.3. Creating an installation source on an NFS server

This procedure describes how to place the installation source on an NFS server. This method allows you to install multiple systems from a single source, without having to connect to physical media. A network-based installation is convenient when used with a TFTP server, as it allows you to boot the installation from the network. This approach eliminates the need to create physical media and simultaneously deploys Red Hat Enterprise Linux on multiple systems.

Prerequisites

- You have downloaded a Binary DVD image. See [Section 4.5, “Download the installation ISO image”](#) for more information.
- You have created a bootable CD, DVD, or USB device from the image file. See [Section 4.6, “Create installation media”](#) for more information.
- You have verified that your firewall allows the server you are installing to access the remote installation source. The following table lists the ports that must be open for each type of network-based installation. See the *Red Hat Enterprise Linux Security Guide* for more information.

Table 4.2. Ports for network-based installation

Protocol used	Ports to open
FTP	21
HTTP	80
HTTPS	443
NFS	2049, 111, 20048
TFTP	69

Procedure



NOTE

The NFS installation method uses the Binary DVD ISO image in a Network File System server's exported directory, which the system must be able to read. To perform an NFS-based installation, another system must act as the NFS host. This procedure is a basic outline of the process. The steps to set up an NFS server vary depending on the system's architecture, operating system, package manager, and service manager.

1. Install the `nfs-utils` package by running the following command as root:

```
# yum install nfs-utils
```

2. Copy the Binary DVD ISO image to a directory on the NFS server.
3. Open the `/etc/exports` file using a text editor and add a line with the following syntax:

```
/exported_directory/ clients
```

4. Replace `/exported_directory/` with the full path to the directory holding the ISO image. Instead of `clients`, use the host name or IP address of the computer that is to be installed from this NFS server, the subnetwork from which all computers are to have access the ISO image, or the asterisk sign (*) if you want to allow any computer with network access to the NFS server to use the ISO image. See the `exports(5)` man page for detailed information about the format of this field.

A basic configuration that makes the `/rhel8-install/` directory available as read-only to all clients is:

```
/rhel8-install *
```

5. Save the `/etc/exports` file and exit the text editor.
6. Start the `nfs` service:

```
# systemctl start nfs.service
```

If the service was running before you changed the `/etc/exports` file, enter the following command, in order for the running NFS server to reload its configuration:

```
# systemctl reload nfs.service
```

The ISO image is now accessible over NFS and ready to be used as an installation source.



NOTE

When configuring the installation source, use `nfs:` as the protocol, the server's host name or IP address, the colon sign (`:`), and the directory holding the ISO image. For example, if the server's host name is `myserver.example.com` and you have saved the ISO image in `/rhel8-install/`, specify `nfs:myserver.example.com:/rhel8-install/` as the installation source.

4.7.4. Creating an installation source using HTTP or HTTPS

This procedure describes how to create an installation source for a network-based installation using an installation tree, which is a directory containing extracted contents of the Binary DVD ISO image and a valid `.treeinfo` file. The installation source is accessed over HTTP or HTTPS.

Prerequisites

- You have downloaded a Binary DVD image. See [Section 4.5, “Download the installation ISO image”](#) for more information.
- You have created a bootable CD, DVD, or USB device from the image file. See [Section 4.6, “Create installation media”](#) for more information.

Procedure

1. Install the **httpd** package by running the following command as root:

```
# yum install httpd
```



WARNING

If your Apache web server configuration enables SSL security, verify that you only enable the TLSv1 protocol, and disable SSLv2 and SSLv3. This is due to the POODLE SSL vulnerability (CVE-2014-3566). See <https://access.redhat.com/solutions/1232413> for details.



IMPORTANT

If you use a HTTPS server with a self-signed certificate, you must boot the installation program with the **noverifyssl** option.

2. Copy the Binary DVD ISO image to the HTTP(S) server.

3. Mount the Binary DVD ISO image, using the **mount** command, to a suitable directory:

```
# mount -o loop,ro -t iso9660 /image_directory/image.iso
/mount_point/
```

4. Replace `/image_directory/image.iso` with the path to the Binary DVD ISO image.
5. Replace `/mount_point/` with the path to the directory where you want to locate the contents of the ISO image.
6. Copy the files from the mounted image to the HTTP server root. This command creates the **/var/www/html/rhel8-install/directory** with the contents of the image.

```
# cp -r /mnt/rhel8-install/ /var/www/html/
```

7. Start the httpd service:

```
# systemctl start httpd.service
```

The installation tree is now accessible and ready to be used as the installation source.



NOTE

When configuring the installation source, use `http://` or `https://` as the protocol, the server's host name or IP address, and the directory in which you have stored the files from the ISO image, relative to the HTTP server root. For example, if you are using HTTP, the server's host name is `myserver.example.com`, and you have copied the files from the image to `/var/www/html/rhel8-install/`, specify <http://myserver.example.com/rhel8-install/> as the installation source.

Additional resources

- For more information about HTTP and FTP servers, see the *Red Hat Enterprise Linux System Administrator's Guide*.

4.7.5. Creating an installation source using FTP

This procedure describes how to create an installation source for a network-based installation using an installation tree, which is a directory containing extracted contents of the Binary DVD ISO image and a valid `.treeinfo` file. The installation source is accessed over FTP.

Prerequisites

- You have downloaded a Binary DVD image. See [Section 4.5, “Download the installation ISO image”](#) for more information.
- You have created a bootable CD, DVD, or USB device from the image file. See [Section 4.6, “Create installation media”](#) for more information.

Procedure

1. Install the vsftpd package by running the following command as root:

```
# yum install vsftpd
```

2. Optionally, open and edit the `/etc/vsftpd/vsftpd.conf` configuration file in a text editor. For available options, see the **vsftpd.conf(5)** man page. This procedure assumes that default options are used.



WARNING

If you configured SSL/TLS security in your `vsftpd.conf` file, make sure to only enable the TLSv1 protocol, and disable SSLv2 and SSLv3. This is due to the POODLE SSL vulnerability (CVE-2014-3566). See <https://access.redhat.com/solutions/1234773> for details.

3. Copy the Binary DVD ISO image to the FTP server.
4. Mount the Binary DVD ISO image, using the `mount` command, to a suitable directory:

```
# mount -o loop,ro -t iso9660 /image_directory/image.iso
/mount_point
```

5. Replace `/image_directory/image.iso` with the path to the Binary DVD ISO image
6. Replace `/mount_point` with the path to the directory where you want to locate the contents of the ISO image.
7. Copy the files from the mounted image to the FTP server root:

```
# cp -r /mnt/rhel8-install/ /var/ftp/
```

This command creates the `/var/ftp/rhel8-install/` directory with the content of the image.

8. Start the `vsftpd` service:

```
# systemctl start vsftpd.service
```

If the service was running before you changed the `/etc/vsftpd/vsftpd.conf` file, restart it to load the edited file.

9. To restart, run the following command:

```
# systemctl restart vsftpd.service
```

The installation tree is now accessible and ready to be used as the installation source.

**NOTE**

When configuring the installation source, use `ftp://` as the protocol, the server's host name or IP address, and the directory in which you have stored the files from the ISO image, relative to the FTP server root. For example, if the server's host name is `myserver.example.com` and you have copied the files from the image to `/var/ftp/rhel8-install/`, specify <ftp://myserver.example.com/rhel8-install/> as the installation source.

CHAPTER 5. BOOTING THE INSTALLATION

The following section explains how to boot the installation program on AMD64, Intel 64, and ARM 64 architectures.

5.1. BOOTING THE INSTALLATION FROM A USB, CD, OR DVD

After you have made a bootable USB flash drive, or a CD or DVD, you are ready to boot the installation. The following steps are generic. Consult your hardware manufacturer's documentation for specific instructions for your system.

Prerequisite

You successfully started the installation program from a bootable device and the boot menu is displayed.

Procedure

1. The initial installer screen presents you with several options defined by sets of **boot options** - commands that tell the system how to proceed. Select the option you want to use by highlighting an entry using the arrow keys.
2. Optional: Press **E** to change the boot options. The menu screen enters edit mode and you can change the predefined command line. Add or remove boot options as needed.
For example, if you want to load a Kickstart file during the installation, add the **inst.ks=** option and enter the location of the Kickstart file immediately after the equals sign.
3. Press **Enter** to confirm your choice.

The **Welcome to Red Hat Enterprise Linux** window opens.

Additional Resources

- [Chapter 6, *Installing using the Graphical User Interface*](#) provides information on installing Red Hat Enterprise Linux using the Graphical User Interface.
- [Section 5.3, “Boot options reference”](#) provides a list of available boot options you can use on the boot command line.
- [<Add in link to advanced install for Kickstart>](#) describes what a Kickstart file is and how to create one.

5.2. BOOTING THE INSTALLATION FROM A NETWORK USING PXE

This procedure describes how to boot from a network using PXE.

Prerequisites

- You have configured a TFTP server, and there is a network interface in your system that supports PXE.
- You have configured your system to boot from the network interface. This option is in the BIOS, and can be labeled **Network Boot** or **Boot Services**.
- You have verified that the BIOS is configured to boot from the specified network interface.

Some BIOS systems specify the network interface as a possible boot device, but do not support the PXE standard. See your hardware's documentation for more information. Once you have properly enabled PXE booting, the system can boot the Red Hat Enterprise Linux installation program without any other media.

Procedure



NOTE

This procedure requires the use of a physical network connection, for example Ethernet. It does work with a wireless connection.

1. Verify that the network cable is attached. The link indicator light on the network socket should be lit, even if the computer is not switched on.
2. Switch on the system.
Depending on your hardware, some network setup and diagnostic information can be displayed before your system connects to a PXE server. Once connected, a menu is displayed according to the PXE server configuration.
3. Press the number key that corresponds to the desired option. The installation program starts, and the boot menu is displayed, containing a variety of boot options.

5.3. BOOT OPTIONS REFERENCE

The installation program boot option reference content is available at [upstream version](#).

CHAPTER 6. INSTALLING USING THE GRAPHICAL USER INTERFACE

This section describes how to install Red Hat Enterprise Linux using the Graphical User Interface (GUI). The GUI is the preferred method of installing Red Hat Enterprise Linux when you boot the system from a CD, DVD, or USB flash drive.



NOTE

- The instructions provided in this section are for AMD and Intel 64-bit, 64-bit ARM, IBM Power Systems, and IBM Z architectures.
- There may be some variance between the online help and the content that is published on the Customer Portal. For the latest updates, see the installation content on the Customer Portal.

6.1. GRAPHICAL INSTALLATION WORKFLOW

You can customize the Red Hat Enterprise Linux graphical installation configuration settings for your specific requirements.

1. Configure your language, location, and localization settings.
2. Select your installation source and required software packages.
3. Configure system memory, security, system purpose, and disk partitions.
4. Configure storage.
5. Configure the network.
6. Create a user account and password, and begin the installation.
7. Complete the installation.



NOTE

When installing from a network location, you must configure the network before you can select the packages you want to install.

6.2. CONFIGURING LANGUAGE AND LOCATION SETTINGS

The selected language is used during the installation program and on the installed system.

Prerequisites

1. You created installation media. See [Section 4.6, “Create installation media”](#) for more information.
2. You booted the installation. See [Chapter 5, *Booting the installation*](#) for more information.

Procedure


1. From the left-hand pane of the **Welcome to Red Hat Enterprise Linux** window, select a language of your choice. Alternatively, type your preferred language into the **Search** field.



NOTE

A language is pre-selected by default. If network access is configured, that is, if you booted from a network server instead of local media, the pre-selected language is determined by the automatic location detection feature of the **GeolP** module. If you used the **inst.lang=** option on the boot command line or in your PXE server configuration, then this language is selected.

2. From the right-hand pane of the **Welcome to Red Hat Enterprise Linux** window, select a location specific to your region.
3. Click **Continue** to proceed to the [Section 6.3, “Installation summary”](#) window.



IMPORTANT

- If you are installing a pre-release version of Red Hat Enterprise Linux, a warning message is displayed about the pre-release status of the installation media. Click **I want to proceed** to continue with the installation, or **I want to exit** to quit the installation and reboot the system.

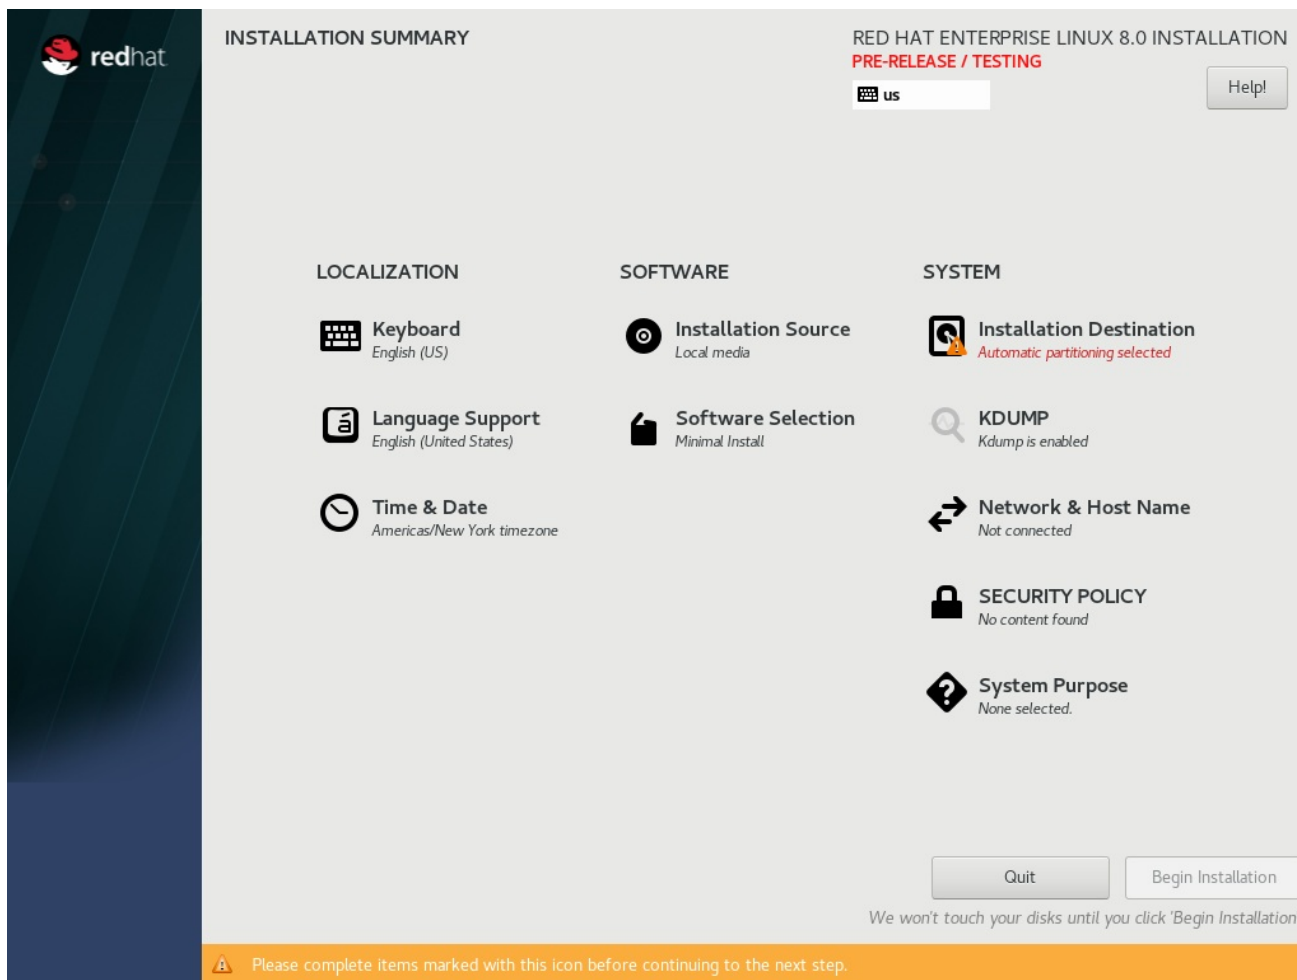
Additional resources

For information on how to change language and location settings during the installation program, see: [Section 6.4, “Localization settings”](#)

6.3. INSTALLATION SUMMARY

The **Installation Summary** window is the central location for the Red Hat Enterprise Linux 8 installation program.

Figure 6.1. Installation summary

**NOTE**

If you used a Kickstart option or a boot option to specify an installation repository on a network, but no network is available at the start of the installation, the installation program displays the **Network Configuration** window to set up a network connection prior to displaying the [Installation Summary window.

The **Installation Summary** window contains three categories:

- **LOCALIZATION** enables you to configure Keyboard, Language Support, and Time and Date.
- **SOFTWARE** enables you to configure Installation Source and Software Selection.
- **SYSTEM** enables you to configure Installation Destination, KDUMP, Network and Host Name, Security Policy, and System Purpose.

A category can have a different status depending on where it is in the installation program.

Table 6.1. Installation Summary category status

Category status	Status	Description
-----------------	--------	-------------

Category status	Status	Description
Warning symbol type 1	Yellow triangle with an exclamation mark and red text	Requires attention before installation. For example, Installation Destination requires attention as you have to confirm the default automatic partitioning variant.
Warning symbol type 2	Greyed out and with a warning symbol (yellow triangle with an exclamation mark)	The installation program is configuring a category and you must wait for it to finish before accessing the window.

**NOTE**

A warning message is displayed at the bottom of the **Installation Summary** window and the **Begin Installation** button is disabled until you configure all of the required categories.

Additional resources

- For information on how to configure Localization settings, see: [Section 6.4, “Localization settings”](#)
- For information on how to configure Software settings, see: [Section 6.5, “Software settings”](#)
- For information on how to configure System settings, see: [Section 6.6, “System Settings”](#)

6.4. LOCALIZATION SETTINGS

This section describes how to configure your keyboard, language support, and time and date settings. It does not detail all aspects of the GUI, only those that are required to complete the task.

**IMPORTANT**

If you use a layout that cannot accept Latin characters, such as **Russian**, you are advised to also add the **English (United States)** layout and configure a keyboard combination to switch between the two layouts. If you only select a layout that does not have Latin characters, you may be unable to enter a valid **root** password and user credentials later in the installation process. This can prevent you from completing the installation.

6.4.1. Configuring keyboard, language, and time and date settings

**NOTE**

Keyboard, Language, and Time and Date Settings were configured as part of [Section 6.2, “Configuring language and location settings”](#). If you want to change any of the settings complete the following steps, otherwise proceed to [Section 6.5, “Software settings”](#).

Procedure

1. From the **Installation Summary** window, click **Keyboard**. The default layout depends on the option selected in [Section 6.2, “Configuring language and location settings”](#).
 - a. Click **+** to open the **Add a Keyboard Layout** window and change to a different layout.
 - b. Select a layout by browsing the list or use the **Search** field.
 - c. Select the required layout and click **Add**. The new layout appears under the default layout.
 - d. Click **Options** to optionally configure a keyboard switch that can be used to cycle between available layouts. The **Layout Switching Options** window opens.
 - e. To configure key combinations for switching, select one or more key combinations and click **OK** to confirm your selection.



NOTE

Once you select a layout, clicking the **Keyboard** button opens a new dialog box that displays a visual representation of the selected layout.

- a. Click **Done** to apply the settings and return to [Section 6.3, “Installation summary”](#).
2. From the **Installation Summary** window, click **Language Support**. The **Language Support** window opens. The left pane lists the available language groups. If at least one language from a group is configured, a check mark is displayed and the supported language is highlighted.
 - a. From the left pane, click a group to select additional languages, and from the right pane, select regional options. Repeat this process for all languages.
 - b. Click **Done** to apply the changes and return to [Section 6.3, “Installation summary”](#).
3. From the **Installation Summary** window, click **Time & Date**. The **Time & Date** window opens.



NOTE

The **Time & Date** window is automatically configured based on the settings you selected in [Section 6.2, “Configuring language and location settings”](#).

The list of cities and regions come from the Time Zone Database (**tzdata**) public domain that is maintained by the Internet Assigned Numbers Authority (IANA). Red Hat can not add cities or regions to this database. You can find more information at the [IANA official website](#).

- a. From the **Region** drop-down menu, select a region.



NOTE

You can select a time zone relative to Greenwich Mean Time (GMT) without setting your location to a specific region. To do so, select **Etc** as your region.

- b. From the **City** drop-down menu, select the city, or the city closest to your location in the same time zone.

- c. Toggle the **Network Time** switch to enable or disable network time synchronization using the Network Time Protocol (NTP).

**NOTE**

Enabling the Network Time switch keeps your system time correct as long as the system can access the internet. By default, one NTP pool is configured; you can add a new option, or disable or remove the default options by clicking the gear wheel button next to the **Network Time** switch.

- d. Click **Done** to apply the changes and return to [Section 6.3, “Installation summary”](#).

**NOTE**

If you disable network time synchronization, the controls at the bottom of the window become active, allowing you to set the time and date manually.

6.5. SOFTWARE SETTINGS

This section describes how to configure your installation source and software selection settings. It does not detail all aspects of the GUI, only those that are required to complete the task.

6.5.1. Configuring installation source

This section details how to install and configure the full installation image, which is the **recommended** method of installing Red Hat Enterprise Linux 8.

Prerequisites

- You have downloaded the full installation image as detailed in [Section 4.5, “Download the installation ISO image”](#).
- You have created a bootable physical media as detailed in [Section 4.6.2, “Creating a bootable USB device on Linux”](#).
- The **Installation Summary** window is open.

**NOTE**

When the **Installation Summary** window first opens, the installation program attempts to configure an installation source based on the type of media that was used to boot the system. The full Red Hat Enterprise Linux Server DVD configures the source as local media.

Procedure

1. From the **Installation Summary** window, click **Installation Source**. The **Installation Source** window opens.
 - a. Review the **Auto-detected installation** section to verify the details. This option is selected by default if you started the installation program from media containing an installation source, for example, a DVD.
 - b. Click **Verify** to check the media integrity.

- c. Review the **Additional repositories** section and note that the **Appstream** checkbox is selected by default.



IMPORTANT

- No additional configuration is necessary as the BaseOS and Appstream repositories are installed as part of the full installation image.
- Do not disable the Appstream repository check box if you want a full Red Hat Enterprise Linux 8 installation.

2. Select the **On the network** option to download and install packages from a network location instead of local media.



NOTE

- This option is only available when a network connection is active. See [Section 6.6.3, “Network and host name”](#) for information on how to set up network connections in the GUI.
- If you do not want to download and install additional repositories from a network location, proceed to [Section 6.5.2, “Configuring software selection”](#).

- a. Select the **On the network** drop-down menu to specify the protocol for downloading packages. This setting depends on the server you want to use.



WARNING

The Appstream repository check box is disabled if you select **On the network** and then decide to revert to **Auto-detected installation**. You must select the Appstream check box to enable the Appstream repository.

- b. Type the server address (without the protocol) into the address field. If you choose NFS, a second input field opens where you can specify custom **NFS mount options**. This field accepts options listed in the **nfs(5)** man page.



IMPORTANT


When selecting an NFS installation source, you must specify the address with a colon (:) character separating the host name from the path. For example:

```
server.example.com:/path/to/directory
```

**NOTE**

The following steps are optional and are only required if a proxy is used for network access.

- c. Click **Proxy setup...** to configure a proxy for a HTTP or HTTPS source.
- d. Select the **Enable HTTP proxy** check box and type the URL into the **Proxy Host** field.
- e. Select the **Use Authentication** check box if the proxy server requires authentication.
- f. Type in your user name and password.
- g. Click **OK** to finish the configuration and exit the **Proxy Setup...** dialog box.

**NOTE**

If your HTTP or HTTPS URL refers to a repository mirror list, select the required option from the **URL type** drop-down list. All environments and add-ons are available for selection once you finish configuring the sources.

3. Click **+** to add a repository.
4. Click **-** to delete a repository.
5. Click the arrow icon to replace the current entries with those that were present at the time you opened the **Installation Source** window.
6. To activate or deactivate a repository, click the check box in the **Enabled** column at each entry in the list.

**NOTE**

You can name your additional repository and configure it the same way as the primary repository on the network.

7. Click **Done** to apply the settings and return to [Section 6.3, “Installation summary”](#).

6.5.2. Configuring software selection

Use the **Software Selection** window to select the required software packages. The packages are organized by Base Environments and Add-Ons.

- **Base environments** are pre-defined packages. Only one base environment can be selected, and availability is dependent on the installation ISO image used as the installation source.
- **Add-Ons** are additional packages for the base environment. You can select multiple add-ons.

The pre-defined environments and add-ons allow you to customize your system, but in a manual installation, you cannot select individual packages to install. To view the packages contained in a specific environment or add-on, see the *repodata/*-comps-variant.architecture.xml* file on the Red Hat Enterprise Linux installation DVD. The XML file describes the packages installed as part of a base environment or add-on. Available environments are marked by the `<environment>` tag, and add-ons are marked by the `<group>` tag.

If you are unsure about which packages to install, Red Hat recommends you to select the **Minimal Install** base environment. Minimal install installs a basic version of Red Hat Enterprise Linux with only a minimal amount of additional software. After the system finishes installing and you log in for the first time, you can use the **Yum package manager** to install additional software. For more information on Yum package manager, see the *Configuring basic system settings* guide.



NOTE

- The **yum group list** command lists all package groups from yum repositories. See the *Configuring basic system settings* guide for more information.
- If you need to control which packages are installed, you can use a Kickstart file and define the packages in the **%packages** section. See *Installing RHEL as an experienced user* for information on Kickstart installations.

Prerequisites

- You have configured the installation source.
- The installation program downloaded a package metadata.
- The Installation Summary window is open.

Procedure

1. From the **Installation Summary** window, click **Software Selection**. The **Software Selection** window opens.
2. From the **Base Environment** pane, select a base environment. Only one environment can be selected.
3. From the **Add-Ons for Selected Environment** pane, select one or more add-ons.



NOTE

If you require a GNOME desktop environment, select the **Graphical Administration Tools** add-on.

4. Click **Done** to apply the settings and return to [Section 6.3, “Installation summary”](#).

6.6. SYSTEM SETTINGS

This section describes how to configure Installation Destination, KDUMP, Network and Host Name, Security Policy, and System Purpose. It does not detail all aspects of the GUI, only those that are required to complete the tasks.

6.6.1. Configuring installation destination

Use the **Installation Destination** window to configure the storage options, for example, the disks used as the installation target for your Red Hat Enterprise Linux installation. At least one disk must be selected. For information about the theory and concepts behind disk partitioning in Linux, see [<ADD LINK TO Partitioning Reference>](#)



WARNING

Back up your data if you plan to use a disk that already contains data. For example, if you want to shrink an existing Microsoft Windows partition and install Red Hat Enterprise Linux 8 as a second system, or if you are upgrading a previous release of Red Hat Enterprise Linux. Manipulating partitions always carries a risk. For example, if the process is interrupted or fails for any reason data on the disk can be lost.

IMPORTANT

Special cases

- Some BIOS types do not support booting from a RAID card. In these instances, the **/boot** partition must be created on a partition outside of the RAID array, such as on a separate hard drive. It is necessary to use an internal hard drive for partition creation with problematic RAID cards. A **/boot** partition is also necessary for software RAID setups. If you have chosen to automatically partition your system, you should manually edit your **/boot** partition.
- To configure the Red Hat Enterprise Linux 8 boot loader to *chain load* from a different boot loader, you must specify the boot drive manually by clicking the **Full disk summary and bootloader** link from the **Installation Destination** window.
- When you install Red Hat Enterprise Linux 8 on a system with both multipath and non-multipath storage devices, the automatic partitioning layout in the installation program creates volume groups that contain a mix of multipath and non-multipath devices. This defeats the purpose of multipath storage. It is recommended that you select either multipath or non-multipath devices on the **Installation Destination** window. Alternatively, proceed to manual partitioning.

Prerequisites

- If you selected the quick installation path, as shown in [Section 6.1, “Graphical installation workflow”](#), you have completed the steps in [Section 6.2, “Configuring language and location settings”](#).
- If you selected the customized installation path, as shown in [Section 6.1, “Graphical installation workflow”](#), you have completed the steps in [Section 6.5, “Software settings”](#).
- The **Installation Summary** window is open.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens.
 - a. **Quick installation path:** Accept the defaults and click **Done** to return to the **Installation Summary** window. The tooltip under **Installation Destination** displays that Automatic Partitioning has been selected.

- i. Click **Begin Installation**.
 - ii. Complete the procedures in [Section 6.9, “Starting the installation program”](#).
- b. **Customized installation path:** Complete the remaining steps in this procedure.
2. From the **Local Standard Disks** section, select the storage device. All storage devices that are used to install Red Hat Enterprise Linux 8 are denoted by a white check mark. Disks that do not have a white check mark are not to be used during the installation, they are ignored if you choose automatic partitioning and they are not available in manual partitioning.

**NOTE**


All locally available storage devices (SATA, IDE and SCSI hard drives, and USB flash drives) are displayed in the **Local Standard Disks**. Any storage devices connected after the installation program has started are not detected.

3. Click **Refresh** if you want to configure additional local storage devices to connect new hard drives.
4. The **Rescan Disks** dialog box opens.

**NOTE**

All storage changes made during the installation program are lost once you click **Rescan Disks**.

- a. Click **Rescan Disks** and wait until the scanning process completes.
 - b. Click **OK** to return to the **Installation Destination** window. All detected disks including any new ones are displayed in the **Local Standard Disks** section.
5. Select one or more storage device(s) that you want to make available during the installation.

**IMPORTANT**

USB storage devices such as flash drives and external disks are also listed under the **Local Standard Disks** section and are available for selection the same way internal hard drives are. If you use a removable drive to install Red Hat Enterprise Linux 8, your system is unusable if you remove the device.

6. To add a specialized storage device, click **Add a disk...** .
The **Storage Device Selection** window opens and lists all storage devices that the installation program has access to. See [Section 6.7.3, “Advanced storage options”](#) for information on adding a specialized disk.
7. Under **Storage Configuration**, select the **Automatic** radio button.

**IMPORTANT**

Automatic partitioning is the **recommended** method of partitioning your storage. You can configure custom partitioning, for more details see [Section 6.8, “Manual partitioning”](#)

8. To reclaim space from an existing partitioning layout, select the **I would like to make additional space available** check box. For example, if a disk you want to use already contains a different operating system and you want to make this system's partitions smaller to allow more room for Red Hat Enterprise Linux 8.
9. Select **Encrypt my data** to encrypt all partitions except the ones needed to boot the system (such as **/boot**) using *Linux Unified Key Setup* (LUKS). Encrypting your hard drive is recommended.
10. Click the **Full disk summary and bootloader** link to select which storage device contains the boot loader. For more information, see [Section 6.6.1.1, “Configuring boot loader”](#).



NOTE

In most cases it is sufficient to leave the boot loader in the default location. Some configurations, for example, systems that require chain loading from another boot loader require the boot drive to be specified manually.

11. Click **Done**.
 - a. If you selected **Encrypt my data** the **Disk Encryption Passphrase** dialog box opens.
 - i. Type your passphrase in the **Passphrase`** and ***Confirm** fields.
 - ii. Click **Save Passphrase** to complete disk encryption.



WARNING

If you lose the LUKS passphrase, any encrypted partitions and the data on them will become completely inaccessible. There is no way to recover a lost passphrase. However, if you perform a Kickstart installation, you can save encryption passphrases and create backup encryption passphrases during the installation.


- b. If you selected automatic partitioning and the **I would like to make additional space available**, or if there is not enough free space on your selected hard drives to install Red Hat Enterprise Linux 8, the **Reclaim Disk Space** dialog box opens and lists all configured disk devices and all partitions on those devices. The dialog box displays information about how much space the system needs for a minimal installation and how much space you have reclaimed.



WARNING

If you used the **Reclaim Space** dialog to **delete** a partition, all data on that partition is lost. If you want to preserve your data, use the **Shrink** option, not the **Delete** option.

- c. Review the displayed list of available storage devices. The **Reclaimable Space** column shows how much space can be reclaimed from each entry.
- d. To reclaim space, select a disk or partition, and click either the **Delete** button to delete that partition, or all partitions on a selected disk, or **Shrink** to use free space on a partition while preserving existing data.



NOTE

Alternatively, you can click **Delete all**, this deletes all existing partitions on all disks and makes this space available to Red Hat Enterprise Linux 8. Existing data on all disks is lost.

12. Click **Reclaim space** to apply the changes and return to [Section 6.3, “Installation summary”](#).



IMPORTANT

No changes to any disks are made until you click **Begin Installation** on the **Installation Summary** window. The **Reclaim Space** dialog only marks partitions for resizing or deletion, but no such action is performed immediately.

6.6.1.1. Configuring boot loader

Red Hat Enterprise Linux 8 uses GRand Unified Bootloader version 2 (**GRUB2**) as its boot loader for AMD64 and Intel 64, IBM Power Systems, and ARM. For IBM Z, the **zipl** boot loader is used.

The boot loader is the first program that runs when the system starts and it is responsible for loading and transferring control to an operating system. **GRUB2** can boot any compatible operating system (including Microsoft Windows) and can also use chain loading to transfer control to other boot loaders for unsupported operating systems.



WARNING

Installing **GRUB2** may overwrite your existing boot loader.

If an operating system is already installed, the Red Hat Enterprise Linux 8 installation program attempts to automatically detect and configure the boot loader to start them. If they are not detected, you can manually configure any additional operating systems after you finish the installation.

If you are installing a Red Hat Enterprise Linux system with more than one disk, you may want to manually specify where the boot loader should be installed.

Procedure

1. From the **Installation Destination** window, click the **Full disk summary and bootloader** link. The **Selected Disks** dialog box opens.
The boot loader is installed on the device of your choice, or on a UEFI system; the **EFI system partition** is created on that device during guided partitioning.

2. To change the boot device, select a device from the list and click **Set as Boot Device**. Only one device can be set as the boot device.
3. To disable a new boot loader installation, select the device currently marked for boot and click **Do not install boot loader**. This ensures **GRUB2** is not installed on any device.



WARNING

If you choose not to install a boot loader, you cannot boot the system directly and you must use another boot method, such as a stand-alone commercial boot loader application. Use this option only if you have another way to boot your system.

The boot loader may also require a special partition to be created, depending on if your system uses BIOS or UEFI firmware, or if the boot drive has a *GUID Partition Table* (GPT) or a **Master Boot Record** (MBR, also known as msdos) label. If you use automatic partitioning, the installation program creates the partition.

6.6.2. Configuring Kdump

Kdump is a kernel crash-dumping mechanism. In the event of a system crash, it captures the contents of the system memory at the moment of failure. This captured memory can be analyzed to find the cause of the crash. If **Kdump** is enabled, it must have a small portion of the system's memory (RAM) reserved to itself. This reserved memory is not accessible to the main kernel.

Procedure

1. From the **Installation Summary** window, click **Kdump**. The **Kdump** window opens.
2. Select the **Enabled kdump** check box.
3. Select either the **Automatic** or **Manual** memory reservation setting.
 - a. If you select **Manual** enter the amount of memory (in megabytes) to be reserved in the **Memory to be reserved** field using the **+** and **-** buttons. The **Usable System Memory** readout below the reservation input field shows how much memory is accessible to your main system once your selected amount of RAM is reserved.
4. Click **Done** to apply the settings and return to [Section 6.3, “Installation summary”](#).



NOTE

The amount of memory you reserve is determined by your system's architecture (AMD64 and Intel 64 have different requirements than IBM Power) as well as the total amount of system memory. In most cases, automatic reservation is satisfactory.



IMPORTANT

Additional settings, such as the location where kernel crash dumps will be saved, can only be configured after the installation using either the **system-config-kdump** graphical interface, or manually in the `/etc/kdump.conf` configuration file.

6.6.3. Network and host name

The **Network and Host name** window is used to configure network interfaces. Options selected here are available both during the installation for tasks such as downloading packages from a remote location, and on the installed system.

6.6.3.1. Configuring network and host name

This section provides information on configuring network and host name.

Procedure


1. From the **Installation Summary** window, click **Network and Host Name**.
2. From the list in the left-hand pane, select an interface. The details are displayed in the right-hand pane.
3. Toggle the **ON/OFF** switch to enable or disable the selected interface.



NOTE

Locally accessible interfaces are automatically detected by the installation program and cannot be manually added or deleted.

4. Click **+** to add a virtual network interface, which can be either: Team, Bond, Bridge, or VLAN.
5. Click **-** to remove a virtual interface.
6. Click **Configure** to change settings such as IP addresses, DNS servers, or routing configuration for an existing interface (both virtual and physical).
7. Type a host name for your system in the **Host Name** field.



NOTE

- There are several types of network device naming standards used to identify network devices with persistent names such as **em1** or **wl3sp0**.
- The host name can be either a fully-qualified domain name (FQDN) in the format `hostname.domainname`, or a short host name with no domain name. Many networks have a Dynamic Host Configuration Protocol (DHCP) service that automatically supplies connected systems with a domain name. To allow the DHCP service to assign the domain name to this machine, only specify the short host name. The value **localhost.localdomain** means that no specific static host name for the target system is configured, and the actual host name of the installed system is configured during the processing of the network configuration, for example, by NetworkManager using DHCP or DNS.

8. Click **Apply** to apply the host name to the installer environment.

6.6.3.2. Adding a virtual network interface

Procedure

1. From the **Network & Host name** window, click the **+** button to add a virtual network interface. The **Add a device** dialog opens.
2. Select one of the four available types of virtual interfaces:
 - a. **Bond**: NIC (*Network Interface Controller*) Bonding, a method to bind multiple physical network interfaces together into a single bonded channel.
 - b. **Bridge**: Represents NIC Bridging, a method to connect multiple separate networks into one aggregate network.
 - c. **Team**: NIC Teaming, a new implementation to aggregate links, designed to provide a small kernel driver to implement the fast handling of packet flows, and various applications to do everything else in user space.
 - d. **Vlan** (*Virtual LAM*): A method to create multiple distinct broadcast domains which are mutually isolated.
3. Select the interface type and click **Add**. An editing interface dialog box opens, allowing you to edit any available settings for your chosen interface type. For more information see [Section 6.6.3.3, “Editing network interface configuration”](#).
4. Click **Save** to confirm the virtual interface settings and return to the **Network & Host name** window.



NOTE

- If you need to change the settings of a virtual interface, select it and click **Configure**.

6.6.3.3. Editing network interface configuration

This section contains information on the most important settings for a typical wired connection used during installation. Many of the available options do not have to be changed and are not carried over to the installed system. Configuration of other types of networks is broadly similar, although the specific configuration parameters may be different.



NOTE

On IBM Z, you cannot add a new connection as the network subchannels need to be grouped and set online beforehand, and this is currently only done in the booting phase.

Procedure

1. To configure a network connection manually, select the interface from the **Network and Host name** window and click **Configure**.
An editing dialog specific to the selected interface opens.

**NOTE**

The options presented depend on the connection type - the available options are slightly different depending on whether it is a physical interface (wired or wireless network interface controller) or a virtual interface (Bond, Bridge, Team, or Vlan) that was previously configured in [Section 6.6.3.2, “Adding a virtual network interface”](#).

The three most common and useful options in the editing dialog are:

6.6.3.4. Enabling or Disabling the Interface Connection

Procedure

1. Click the **General** tab.
2. Select the **Automatically connect to this network when it is available** check box to enable connection by default.

**NOTE**

- When enabled on a wired connection, the system typically connects during startup (unless you unplug the network cable). On a wireless connection, the interface attempts to connect to any known wireless networks in range.
- You can enable or disable all users on the system from connecting to this network using the **All users may connect to this network** option. If you disable this option, only **root** will be able to connect to this network.
- It is not possible to only allow a specific user other than **root** to use this interface, as no other users are created at this point during the installation. If you need a connection for a different user, you must configure it after the installation.

3. Click **Save** to apply the changes and return to the **Network and Host name** window.

6.6.3.5. Setting up Static IPv4 or IPv6 Settings

By default, both IPv4 and IPv6 are set to automatic configuration depending on current network settings. This means that addresses such as the local IP address, DNS address, and other settings will be detected automatically when the interface connects to a network. In many cases, this is sufficient, but you can also provide static configuration in the **IPv4 Settings** and **IPv6 Settings** tabs.


Procedure

1. To set static network configuration, navigate to one of IPv Settings tabs and from the **Method** drop-down menu, select a method other than **Automatic**, for example, **Manual**. The **Addresses** pane is enabled.

**NOTE**

In the **IPv6 Settings** tab, you can also set the method to **Ignore** to disable IPv6 on this interface.

2. Click **Add** and enter your address settings.
3. Type the IP addresses in the **Additional DNS servers** field; it accepts one or more IP addresses of DNS servers, for example, **10.0.0.1, 10.0.0.8**.
4. Select the **Require IPvX addressing for this connection to complete** check box.

**NOTE**

Select this option in the **IPv4 Settings** or **IPv6 Settings** tabs to only allow this connection if IPv4 or IPv6 was successful. If this option remains disabled for both IPv4 and IPv6, the interface is able to connect if configuration succeeds on either IP protocol.

5. Click **Save** to apply the changes and return to the **Network & Host name** window.

6.6.3.6. Configuring Routes

Procedure

1. In the **IPv4 Settings** and **IPv6 Settings** tabs, click **Routes** to configure routing settings for a specific IP protocol on an interface. An editing routes dialog specific to the interface opens.
2. Click **Add** to add a route.
3. Select the **Ignore automatically obtained routes** check box to configure at least one static route and want to disable all routes not specifically configured.
4. Select the **Use this connection only for resources on its network** check box to prevent the connection from becoming the default route.

**NOTE**

This option can be selected even if you did not configure any static routes. This route is used only to access certain resources, such as intranet pages that require a local or VPN connection. Another (default) route is used for publicly available resources. Unlike the additional routes configured, this setting is transferred to the installed system. This option is only useful when more than one interface is configured.

5. Click **OK** to save your settings and return to the editing routes dialog specific to the interface.
6. Click **Save** to apply the settings and return to the **Network and Host Name** window.

6.6.4. Security policy

This section provides information on the Red Hat Enterprise Linux 8 security policy add-on and how to configure it for use on your system.

6.6.4.1. About security policy

The Red Hat Enterprise Linux 8 security policy adheres to restrictions and recommendations (compliance policies) defined by the Security Content Automation Protocol (SCAP) standard. The

packages are automatically installed. However, by default, no policies are enforced and therefore no checks are performed during or after installation unless specifically configured.

Applying a security policy is not a mandatory feature of the installation program. If you apply a security policy to the system, it is installed using restrictions and recommendations defined in the selected profile. The **openscap-scanner** package is added to your package selection, providing a preinstalled tool for compliance and vulnerability scanning. After the installation finishes, the system is automatically scanned to verify compliance. The results of this scan are saved to the **/root/openscap_data** directory on the installed system. You can also load additional profiles from a HTTP, HTTPS, or FTP server.

6.6.4.2. Configuring a security policy

This section provides information on how to configure a security policy.

Prerequisites

- The **Installation Summary** window is open.

Procedure

1. From the **Installation Summary** window, click **Security Policy**. The **Security Policy** window opens.
2. To enable security policies on the system, toggle the **Apply security policy** switch to **ON**.
3. Select one of the profiles listed in the top pane.
4. Click **Select profile**.
Profile changes that you must apply before installation appear in the bottom pane.



NOTE

The default profiles do not require changes before installation. However, loading a custom profile can require pre-installation tasks.

5. Click **Change content** to use a custom profile. A separate window opens allowing you to enter a URL for valid security content.
 - a. Click **Fetch** to retrieve the URL.
 - b. Click **Use SCAP Security Guide** to return to the **Security Policy** window.



NOTE

Custom profiles can be loaded from a **HTTP**, **HTTPS**, or **FTP** server. Use the full address of the content, including the protocol such as **http://**. A network connection must be active before you can load a custom profile. The content type is detected automatically by the installation program.

6. Click **Done** to apply the settings and return to the **Installation Summary** window.

6.6.4.3. Related information

- **scap-security-guide(8)** - The manual page for the **scap-security-guide** project provides information on SCAP security profiles, including examples on how to utilize the provided benchmarks using the OpenSCAP utility.
- Red Hat Enterprise Linux security compliance information is available in the *Configuring and managing security* guide.

6.6.5. System Purpose

System administrators use System Purpose to record the intended use of a Red Hat Enterprise Linux 8 system by the organization. When you set a system's purpose, the entitlement server receives information that helps auto-attach a subscription that satisfies the intended use of the system.

You can enter System Purpose data in the following ways:

- During Composer image creation.
- During installation using the graphical user interface.
- Using Kickstart automation scripts.
- Using Runtime operations, for example, the command line and the web console.

You can configure the following components:

- **Role:**
 - Red Hat Enterprise Linux Server
 - Red Hat Enterprise Linux Workstation
 - Red Hat Enterprise Linux Compute Node
- **Service Level Agreement:**
 - Premium
 - Standard
 - Self-Support
- **Usage:**
 - Production
 - Disaster Recovery
 - Development/Test

Benefits include:

- In-depth system-level information for system administrators and business operations.
- Reduced overhead when determining why a system was procured and its intended purpose.
- Improved customer experience of Subscription Manager auto-attach as well as automated discovery and reconciliation of system usage.

Additional resources

- For more information on Composer, see the *Composing a customized RHEL system image* guide.
- For more information on Kickstart, see the *Performing an advanced RHEL installation* guide.
- For more information on Subscription Manager, see the *Using and Configuring Subscription Manager* guide.

6.6.5.1. Configuring system purpose using the graphical user interface



NOTE

While it is strongly recommended that you configure System Purpose, it is an optional feature of the Red Hat Enterprise Linux 8 installation program. If you want to enable System Purpose after the installation completes, you can do so using the **syspurpose** Kickstart command.

Prerequisite

The **Installation Summary** window is open.

Procedure

1. From the **Installation Summary** window, click **System Purpose**.
2. Select the required system role from the **Role** pane.
3. Select the required service level agreement from the **Red Hat Service Level Agreement** pane.
4. Select the required usage from the **Usage** pane.
5. Click **Done** to apply the settings and return to the **Installation Summary** window.

The System Purpose data is now available for Subscription Manager to auto-attach to the system.

6.6.5.2. Configuring System Purpose using Kickstart

Use the **syspurpose** command to configure System Purpose in a Kickstart configuration file.

The following actions are available:

role

Set the intended role of the system. This action uses the following format:

```
syspurpose --role=
```

The role assigned can be:

- **Red Hat Enterprise Linux Server**
- **Red Hat Enterprise Linux Workstation**
- **Red Hat Enterprise Linux Compute Node**

sla

Set the intended sla of the system. This action uses the following format:

```
syspurpose --sla=
```

The sla assigned can be:

- **Premium**
- **Standard**
- **Self-Support**

usage

Set the intended usage of the system. This action uses the following format:

```
syspurpose --usage=
```

The usage assigned can be:

- **Production**
- **Disaster Recovery**
- **Development/Test**

addon



WARNING

The addon functionality is not available in Red Hat Enterprise Linux 8.

Set additional layered products or features. To add multiple items specify **--addon** multiple times, once per layered product or feature:

```
syspurpose --addon=
```

Additional resources

- For more information on Kickstart, see the *Performing an advanced RHEL installation* guide.
- For more information on Subscription Manager, see the *Using and Configuring Subscription Manager* guide.

6.7. STORAGE DEVICES

You can install Red Hat Enterprise Linux 8 on a large variety of storage devices. You can configure basic,

locally accessible, storage devices in the **Installation Destination** window. Basic storage devices directly connected to the local system, such as hard disk drives and solid-state drives, are displayed in the **Local Standard Disks** section of the window. On System z, this contains activated Direct Access Storage Devices (DASDs).



WARNING

A known issue prevents DASDs configured as HyperPAV aliases from being automatically attached to the system after the installation is finished. These storage devices are available during the installation, but are not immediately accessible after you finish installing and reboot. To attach HyperPAV alias devices, add them manually to the system's `/etc/dasd.conf` configuration file.

6.7.1. Storage device selection

The storage device selection window lists all storage devices to which Anaconda has access. Depending on your system and available hardware, some tabs might not be displayed. The devices are grouped under the following tabs:

Multipath Devices

Storage devices accessible through more than one path, such as through multiple SCSI controllers or Fiber Channel ports on the same system.



IMPORTANT

The installation program only detects multipath storage devices with serial numbers that are 16 or 32 characters long.

Other SAN Devices

Devices available on a Storage Area Network (SAN).

Firmware RAID

Storage devices attached to a firmware RAID controller.

NVDIMM Devices

Under specific circumstances, Red Hat Enterprise Linux 8 can boot and run from (NVDIMM) devices in sector mode on the Intel 64 and AMD64 architectures.

System z Devices

Storage devices, or Logical Units (LUNs), attached through the zSeries Linux FCP (Fiber Channel Protocol) driver.

6.7.2. Filtering storage devices

In the storage device selection window you can filter storage devices either by their World Wide Identifier (WWID) or by the port, target, or logical unit number (LUN).

Procedure


1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk....** The storage devices selection window opens.
3. Click the **Search by** tab to search by port, target, LUN, or WWID. Searching by WWID or LUN requires additional values in the corresponding input text fields.
4. Click the required option from the **Search** drop-down menu.
5. Click **Find** to start the search. Each device is presented on a separate row with a corresponding check box.
6. Select the check box to make the device available during the installation process. Later in the installation process you can choose to install Red Hat Enterprise Linux 8 on any of the selected devices and you can choose to automatically mount any of the other selected devices as part of the installed system.



NOTE

- Selected devices are not automatically erased by the installation process and selecting a device does not put the data stored on the device at risk.
- You can add devices to the system after installation by modifying the `/etc/fstab` file.

7. Click **Done** to return to the **Installation Destination** window.



IMPORTANT

Any storage devices that you do not select are hidden from Anaconda entirely. To chain load the boot loader from a different boot loader, select all the devices presented.

6.7.3. Advanced storage options

To use an advanced storage device, you can configure an iSCSI (SCSI over TCP/IP) target or FCoE (Fibre Channel over Ethernet) SAN (Storage Area Network).

To use iSCSI storage devices for the installation, Anaconda must be able to discover them as iSCSI targets and be able to create an iSCSI session to access them. Each of these steps might require a user name and password for CHAP (Challenge Handshake Authentication Protocol) authentication. Additionally, you can configure an iSCSI target to authenticate the iSCSI initiator on the system to which the target is attached (reverse CHAP), both for discovery and for the session. Used together, CHAP and reverse CHAP are called mutual CHAP or two-way CHAP. Mutual CHAP provides the greatest level of security for iSCSI connections, particularly if the user name and password are different for CHAP authentication and reverse CHAP authentication.



NOTE

Repeat the iSCSI discovery and iSCSI login steps to add all required iSCSI storage. However, you cannot change the name of the iSCSI initiator after you attempt discovery for the first time. To change the iSCSI initiator name, you must restart the installation.

6.7.3.1. Discovering and starting an iSCSI session

This section describes how to discover and start an iSCSI session.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk...** The storage devices selection window opens.
3. Click **Add iSCSI target...** The **Add iSCSI Storage Target** window opens.
4. Enter the IP address of the iSCSI target in the **Target IP Address** field.
5. Type a name in the **iSCSI Initiator Name** field for the iSCSI initiator in iSCSI qualified name (IQN) format. A valid IQN entry contains the following information:
 - The string **iqn.** (note the period).
 - A date code that specifies the year and month in which your organization's Internet domain or subdomain name was registered, represented as four digits for the year, a dash, and two digits for the month, followed by a period. For example, represent September 2010 as **2010-09.**
 - Your organization's Internet domain or subdomain name, presented in reverse order with the top-level domain first. For example, represent the subdomain **storage.example.com** as **com.example.storage.**
 - A colon followed by a string that uniquely identifies this particular iSCSI initiator within your domain or subdomain. For example, **:diskarrays-sn-a8675309.**
A complete IQN is as follows: **iqn.2010-09.storage.example.com:diskarrays-sn-a8675309.** **Anaconda** prepopulates the **iSCSI Initiator Name** field with a name in this format to help you with the structure.

For more information on IQNs, see 3.2.6. *iSCSI Names* in *RFC 3720 - Internet Small Computer Systems Interface (iSCSI)* available from tools.ietf.org and 1. *iSCSI Names and Addresses* in *RFC 3721 - Internet Small Computer Systems Interface (iSCSI) Naming and Discovery* available from tools.ietf.org.
6. Select the **Discovery Authentication Type** drop-down menu to specify the type of authentication to use for iSCSI discovery. The following options are available:
 - No credentials
 - CHAP pair
 - CHAP pair and a reverse pair
7.
 - a. If you selected **CHAP pair** as the authentication type, provide the user name and password for the iSCSI target in the **CHAP Username** and **CHAP Password** fields.
 - b. If you selected **CHAP pair and a reverse pair** as the authentication type, provide the user name and password for the iSCSI target in the **CHAP Username** and **CHAP Password** field and the user name and password for the iSCSI initiator in the **Reverse**

CHAP Username and Reverse CHAP Password fields.

8. Optionally, select the **Bind targets to network interfaces** check box.
9. Click **Start Discovery**.
Anaconda attempts to discover an iSCSI target based on the information provided. If discovery succeeds, the **Add iSCSI Storage Target** window displays a list of all iSCSI nodes discovered on the target.
10. Select the required node check boxes to use for installation.

**NOTE**

The **Node login authentication type** menu provides the same options as the **Discovery Authentication Type** menu. However, if you need credentials for discovery authentication, use the same credentials to log into a discovered node.

11. Click the additional **Use the credentials from discovery** drop-down menu. When the proper credentials have been provided, the **Log In** button becomes available.
12. Click **Log In** to initiate an iSCSI session.

6.7.3.2. Configuring FCoE parameters

This section describes how to configure FCoE parameters.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk...** The storage devices selection window opens.
3. Click **Add FCoE SAN...** . A dialog box opens for you to configure network interfaces for discovering FCoE storage devices.
4. Select a network interface that is connected to an FCoE switch in the **NIC** drop-down menu.
5. Click **Add FCoE disk(s)** to scan the network for SAN devices.
6. Select the required check boxes:
 - Use DCB: *Data Center Bridging* (DCB) is a set of enhancements to the Ethernet protocols designed to increase the efficiency of Ethernet connections in storage networks and clusters. Select the check box to enable or disable the installation program's awareness of DCB. This option should only be enabled for network interfaces that require a host-based DCBX client. Configurations on interfaces that implement a hardware DCBX client should disable the check box.
 - Use auto vlan: *Auto VLAN* indicates whether VLAN discovery should be performed. If this check box is enabled, then the FIP (FCoE Initiation Protocol) VLAN discovery protocol will run on the Ethernet interface once the link configuration has been validated. If they are not

already configured, network interfaces for any discovered FCoE VLANs will be automatically created and FCoE instances will be created on the VLAN interfaces. This option is enabled by default.

7. Discovered FCoE devices are displayed under the **Other SAN Devices** tab in the **Installation Destination** window.

6.7.3.3. Configuring DASD storage devices

This section describes how to configure DASD storage devices.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk....** The storage devices selection window opens.
3. Click **Add DASD**. The **Add DASD Storage Target** dialog box opens and allows you to attach additional DASDs that were not detected when the installation started. The **Add DASD Storage Target** dialog box prompts you to specify a device number, such as **0.0.0204**.
4. Type the device number of the DASD you want to attach in the **Device number** field.
5. Click **Start Discovery**.



NOTE

- If a DASD with the specified device number is found and if it is not already attached, the dialog box closes and the newly-discovered drives appear in the list of drives. You can then select the check boxes to select the drives that should be made available. After you do so, click **Done**. The new DASDs are available for selection (marked as **DASD device 0.0.xxxx**) in the **Local Standard Disks** section of the **Installation Destination** window.
- If you entered an invalid device number, or if the DASD with the specified device number is already attached to the system, an error message appears in the dialog box, explaining the error and prompting you to try again with a different device number.

6.7.3.4. Configuring FCP devices


FCP devices enable IBM Z to use SCSI devices rather than, or in addition to, Direct Access Storage Device (DASD) devices. FCP devices provide a switched fabric topology that enables IBM Z systems to use SCSI LUNs as disk devices in addition to traditional DASD devices.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk....** The storage devices selection window opens.

3. Click **Add ZFCP LUN**. The **Add zFCP Storage Target** dialog box opens allowing you to add a FCP (Fibre Channel Protocol) storage device.
IBM Z requires that any FCP device is entered manually for the installation program to activate FCP LUNs. This can be done either in **Anaconda** interactively, or specified as a unique parameter entry in the parameter or CMS configuration file. The values entered here are unique to each site in which they are set up.
4. Type the 4 digit hexadecimal device number in the **Device number** field.
5. Type the 16 digit hexadecimal World Wide Port Number (WWPN) in the **WWPN** field.
6. Type the 16 digit hexadecimal FCP LUN identifier in the **LUN** field.
7. Click **Start Discovery** to connect to the FCP device.

The newly-added devices are displayed in the **System z Devices** tab of the **Installation Destination** window.



NOTE

- Interactive creation of an FCP device is only possible in graphical mode. It is not possible to interactively configure an FCP device in text mode installation.
- Use only lower-case letters in hex values. If you enter an incorrect value and click **Start Discovery**, the installation program displays a warning. You can edit the configuration information and retry the discovery attempt.
- For more information on these values, consult the hardware documentation and check with your system administrator who set up the network.



IMPORTANT

For an FCP-only installation, remove the **DASD=** from the CMS configuration file or the **rd.dasd=** from the parameter file to indicate that no DASD is present.

6.7.4. Installing to an NVDIMM device

Non-Volatile Dual In-line Memory Module (NVDIMM) devices combine the performance of RAM with disk-like data persistence when no power is supplied. Under specific circumstances, Red Hat Enterprise Linux 8 can boot and run from (NVDIMM) devices.

6.7.4.1. Criteria for using an NVDIMM Device as an installation target

Red Hat Enterprise Linux 8 can be installed to Non-Volatile Dual In-line Memory Module (NVDIMM) devices in sector mode on the Intel 64 and AMD64 architectures, supported by the **nd_pmem** driver.

Using an NVDIMM Device as Storage

To use an NVDIMM device as storage, the following conditions must be satisfied:

- The architecture of the system is Intel 64 or AMD64.
- The NVDIMM device is configured to sector mode. **Anaconda** can reconfigure NVDIMM devices to this mode.

- The NVDIMM device must be supported by the **nd_pmem** driver.

Booting from an NVDIMM Device

Booting from an NVDIMM device is possible under the following conditions:

- All conditions for using the NVDIMM device as storage are satisfied.
- The system uses UEFI.
- The NVDIMM device must be supported by firmware available on the system, or by an UEFI driver. The UEFI driver may be loaded from an option ROM of the device itself.
- The NVDIMM device must be made available under a namespace.

To take advantage of the high performance of NVDIMM devices during booting, place the **/boot** and **/boot/efi** directories on the device. The Execute-in-place (XIP) feature of NVDIMM devices is not supported during booting and the kernel is loaded into conventional memory.

6.7.4.2. Configuring an NVDIMM device using Anaconda

A Non-Volatile Dual In-line Memory Module (NVDIMM) device must be properly configured for use by Red Hat Enterprise Linux 8 and **Anaconda** as an installation target.



WARNING

Reconfiguration of a NVDIMM device process destroys any data stored on the device.

Prerequisites

- A NVDIMM device is present on the system and satisfy all the other conditions for usage as an installation target.
- Anaconda has booted and the **Installation Summary** window is open.

Procedure

1. From the **Installation Summary** window, click **Installation Destination**. The **Installation Destination** window opens, listing all available drives.
2. Under the **Specialized & Network Disks** section, click **Add a disk...** . The storage devices selection window opens.
3. Click the **NVDIMM Devices** tab.
4. To reconfigure a device select it from the list.
If a device is not listed, it is not in sector mode.
5. Click **Reconfigure NVDIMM...** . A reconfiguration dialog opens.

6. Enter the required sector size and click **Start Reconfiguration**.
The supported sector sizes are 512 and 4096 bytes.
7. Once reconfiguration completes click **OK**.
8. Select the device check box.
9. Click **Done** to return to the **Installation Destination** window.
The NVDIMM device you reconfigured is displayed in the **Specialized & Network Disks** section.
10. Click **Done** to return to the **Installation Summary** window.

The NVDIMM device is now available for selection as an installation target. Additionally, if the device meets the requirements for booting, it can be set as such.

6.7.4.3. Configuring an NVDIMM device using kickstart

By default, all Non-Volatile Dual In-line Memory Module (NVDIMM) devices are ignored by the installation program. To manipulate an NVDIMM device in a Kickstart configuration file, use the **nvdimm** command:

```
nvdimm action [options]
```

The following actions are available:

reconfigure

Reconfigure a specific NVDIMM device to a given mode. Additionally, the specified device is implicitly marked in use, so a subsequent **nvdimm use** command for the same device is redundant. This action uses the following format:

```
nvdimm reconfigure [--namespace=NAMESPACE] [--mode=MODE] [--  
sectorsize=SECTORSIZE]
```

- **--namespace=** - The device specification by namespace. For example:

```
nvdimm reconfigure --namespace=namespace0.0 --mode=sector --  
sectorsize=512
```

- **--mode=** - The mode specification. Currently, only the value **sector** is available.
- **--sectorsize=** - Size of a sector for sector mode. For example:

```
nvdimm reconfigure --namespace=namespace0.0 --mode=sector --  
sectorsize=512
```

The supported sector sizes are 512 and 4096 bytes.

use

Specify a NVDIMM device as a target for installation. The device must be already configured in the sector mode. This action uses the following format:

```
nvdimm use [--namespace=NAMESPACE|--blockdevs=DEVICES]
```

- **--namespace=** - Specifies the device by namespace. For example:

```
nvdimm use --namespace=namespace0.0
```

- **--blockdevs=** - Specifies a comma-separated list of block devices corresponding to the NVDIMM devices to be used. The asterisk * wildcard is supported. For example:

```
nvdimm use --blockdevs=pmem0s,pmem1s
nvdimm use --blockdevs=pmem*
```

6.8. MANUAL PARTITIONING

Manual Partitioning allows you to configure your disk partitions and mount points. This defines the file system that Red Hat Enterprise Linux 8 is installed on.



NOTE

Before installation, you should consider whether you want to use partitioned or unpartitioned disk devices. For more information, see the Knowledgebase article at <https://access.redhat.com/solutions/163853>.

An installation of Red Hat Enterprise Linux 8 requires a minimum of one partition but Red Hat recommends using at least the following partitions or volumes: **/**, **/home**, **/boot**, and **swap**. You can also create additional partitions and volumes as you require. See **LINK TO BE ADDED** for more information.



WARNING

It is recommended that you back up data before proceeding. If you are upgrading or creating a dual-boot system, you should back up any data you want to keep on your storage devices. Unforeseen circumstances can result in data loss.

6.8.1. Starting manual partitioning

Prerequisites

- The **Installation Summary** screen is currently displayed.
- All disks are available to the installation program.

Procedure

1. Select disks for installation:
 - a. Click **Installation Destination** to open the **Installation Destination** window.

- b. Select the required disks for installation by clicking the corresponding icon. A selected disk has a check-mark displayed on it.
 - c. Under **Storage Configuration**, select the **Custom** radio-button.
 - d. Optional: To enable storage encryption with LUKS, select the **Encrypt my data** check box.
 - e. Click **Done**.
2. If you selected to encrypt the storage, a dialog box for entering a disk encryption passphrase opens. Type in the passphrase:
- a. Enter the passphrase into the two text fields. To switch keyboard layout, use the keyboard icon.

**WARNING**

In the dialog box for entering the passphrase, you are not able to change keyboard layout. Select the English keyboard layout to enter the passphrase in the installation program.

- b. The dialog box provides an assessment of the passphrase strength. Change your passphrase if necessary.
 - c. Click **Save Passphrase** to save the passphrase.
The Manual Partitioning window opens.
3. Mount points that the installation program has detected are listed in the left-hand pane. The mount points are organized by detected operating system installations. As a result, some file systems may be displayed multiple times if a partition is shared among several installations.
- a. Select the mount points in the left pane; the customizable options are displayed in the right pane.
 - b. If your system contains existing file systems, ensure that enough space is available for the installation. To remove any partitions that you do not want to be present, select them in the list and click the **-** button.
The dialog has a check box to remove all other partitions used by the system to which the deleted partition belongs.
 - c. If there are no existing partitions and you want to create the recommended set of partitions as a starting point, select your preferred partitioning scheme from the left pane (default for Red Hat Enterprise Linux is LVM) and click the **Click here to create them automatically** link.
A **/boot** partition, a **/** (root) volume, and a **swap** volume proportionate to the size of the available storage are created and listed in the left pane. These are the recommended file systems for a typical installation, but you can add additional file systems and mount points.
 - d. Continue with [adding mount points](#), [configuring the individual mount points](#), and [configuring the underlying partitions or volumes](#).

6.8.2. Adding a mount point file system

You can add multiple mount point file systems.

Prerequisites

- Plan for your partitions:
 - To avoid problems with space allocation, first create small partitions with known fixed sizes, such as **/boot**, and then create the remaining partitions, letting the installation program allocate the remaining capacity to them.
 - If you have multiple disks that the system is to reside on, or if they differ in size and a particular partition must be created on the first disk detected by BIOS, then create these partitions first.

Procedure

1. Click **+** to create a new mount point file system. The **Add a New Mount Point** dialog opens.
2. Select one of the preset paths from the **Mount Point** drop-down menu or type your own; for example, select **/** for the root partition or **/boot** for the boot partition.
3. Type the size of the file system in the **Desired Capacity** field; for example, **2GiB**.




WARNING

If you leave the **Desired Capacity** field empty or if you specify a size bigger than available space, then all remaining free space is used.

4. Click **Add mount point** to create the partition and return to the **Manual Partitioning** window.

6.8.3. Configuring a mount point file system

You can set the partitioning scheme for each mount point that was created manually. The available options are **Standard Partition**, **LVM**, and **LVM Thin Provisioning**.



NOTE

- BTRFS support has been deprecated in Red Hat Enterprise Linux 8.
- The **/boot** partition is always located on a standard partition, regardless of the value selected.

Procedure

1. To change the devices that a single non-LVM mount point should be located on, select the required mount point from the left-hand pane.

2. Under the **Device(s)** heading, click **Modify...** . The **Configure Mount Point** dialog opens.
3. Select one or more devices and click **Select** to confirm your selection and return to the **Manual Partitioning** window.
4. Click **Update Settings** to apply the changes.



NOTE

Click the **Rescan** button (circular arrow button) to refresh all local disks and partitions; this is only required after performing advanced partition configuration outside the installation program. Clicking the **Rescan Disks** button resets all configuration changes made in the installation program.

5. In the lower left-hand side of the **Manual Partitioning** window, click the **storage device selected** link to open the **Selected Disks** dialog and review disk information.

6.8.4. Customizing a partition or volume

Customizing a partition or volume is available if you want to set specific settings.

Procedure

1. From the left pane, select the mount point.

Figure 6.2. Customizing Partitions

The screenshot shows the 'MANUAL PARTITIONING' window for 'RED HAT ENTERPRISE LINUX 8.0 INSTALLATION'. The left pane lists partitions: /boot (1024 MiB), / (rhel-root, 17 GiB), and swap (2 GiB). The 'rhel-root' partition is selected. The right pane shows configuration options for 'rhel-root': Mount Point (/), Device(s) (ATA QEMU HARDDISK (sda)), Desired Capacity (17 GiB), Device Type (LVM), File System (xfs), Volume Group (rhel), and Name (root). There are buttons for 'Done', 'Help!', 'Update Settings', and 'Reset All'. A note at the bottom states: 'Note: The settings you make on this screen will not be applied until you click on the main menu's 'Begin Installation' button.'

2. You can customize the following options in the right pane:

- a. Enter the file system's mount point into the **Mount Point** field. For example, if a file system is the root file system, enter `/`; enter `/boot` for the `/boot` file system, and so on. For a swap file system, the mount point should not be set as setting the file system type to **swap** is sufficient.
- b. Enter the size of the file system in the **Desired Capacity** field. You can use common size units such as KiB or GiB. The default is MiB if no other unit is specified.
- c. Select the required device type from the drop-down **Device Type** menu: **Standard Partition**, **LVM**, or **LVM Thin Provisioning**.

**NOTE**

RAID is only available if two or more disks are selected for partitioning. If you choose this type, you can also set the **RAID Level**. Similarly, if you select **LVM**, you can specify the **Volume Group**.

- d. Select the **Encrypt** check box to encrypt the partition or volume. You are required to set a password later in the installation program. The **LUKS Version** drop-down menu is displayed.
- e. Select the required LUKS version from the drop-down menu.
- f. Select the appropriate file system type for this partition or volume from the **File system** drop-down menu.
- g. Select the **Reformat** check box to format an existing partition, or deselect it to retain your data. The newly-created partitions and volumes must be reformatted, and the check box cannot be deselected.
- h. Type a label for the partition in the **Label** field. Labels are used for you to easily recognize and address individual partitions.
- i. Type a name in the **Name** field.

**NOTE**

Note that standard partitions are named automatically when they are created and their name cannot be edited, such as `/boot` being assigned the name **sda1**.

3. Click **Update Settings** to apply your changes and if required, select another partition to customize. Changes are not applied until you start the installation from the **Installation Summary** window.

**NOTE**

Click **Reset All** to discard partition changes and start over.

4. Click **Done** when all file systems and mount points are created and customized. If you chose to encrypt a file system, you are prompted to create a passphrase. A **Summary of Changes** dialog box opens, displaying a summary of all storage actions for the installation program.

5. Click **Cancel & Return to Custom Partitioning** to return to the **Manual Partitioning** window and make additional changes.
6. Click **Accept Changes** to apply the changes and return to the **Installation Summary** window.



IMPORTANT


If **/usr** or **/var** is partitioned separately from the rest of the root volume, the boot process becomes much more complex as these directories contain critical components. In some situations, such as when these directories are placed on an iSCSI drive or an FCoE location, the system is unable to boot, or hangs with a **Device is busy** error when powering off or rebooting.

This limitation only applies to **/usr** or **/var**, not to directories below them. For example, a separate partition for **/var/www** works successfully.

6.8.5. Creating software RAID

This section describes how to create a RAID device. Redundant Arrays of Independent Disks (RAIDs) are constructed from multiple storage devices that are arranged to provide increased performance and, in some configurations, greater fault tolerance.

A RAID device is created in one step and disks are added or removed as necessary. One RAID partition per physical disk is allowed for each device, so the number of disks available to the installation program determines the levels of RAID device available. For example, if your system has two hard drives, the installation program does not allow you to create a RAID10 device, as it requires 4 separate partitions.



NOTE

On IBM Z, the storage subsystem uses RAID transparently. There is no need to set up a software RAID manually.

Prerequisites

- You have selected two or more disks for installation before RAID configuration options are visible. At least two disks are required to create a RAID device.
- You have created a mount point. By configuring a mount point, you configure the RAID device.
- You have selected the **Custom** radio button on the **Installation Destination** window.

Procedure

1. From the left pane of the **Manual Partitioning** window, select the required partition.
2. Under the **Device(s)** section, click **Modify**. The **Configure Mount Point** dialog box opens.
3. Select the disks that are to be included in the RAID device and click **Select**.
4. Click the **Device Type** drop-down menu and select **RAID**.
5. Click the **File System** drop-down menu and select your preferred file system type.

6. Click the **RAID Level** drop-down menu and select your preferred level of RAID.
7. Click **Update Settings** to save your changes.
8. Click **Done** to apply the settings and return to the **Installation Summary** window.

A message is displayed at the bottom of the window if the specified RAID level requires more disks.

6.8.6. Creating an LVM logical volume

Logical Volume Management (LVM) presents a simple logical view of underlying physical storage space, such as hard drives or LUNs. Partitions on physical storage are represented as physical volumes that can be grouped together into volume groups. You can divide each volume group into multiple logical volumes, each of which is analogous to a standard disk partition. Therefore, LVM logical volumes function as partitions that can span multiple physical disks.



NOTE

Note that LVM configuration is only available in the graphical installation program.



IMPORTANT

During text-mode installation, LVM configuration is not available. To create an LVM configuration, press **Ctrl+Alt+F2** to use a different virtual console, and run the **lvm** command. To return to the text-mode installation, press **Ctrl+Alt+F1**.

Procedure

1. From the left-hand pane of the **Manual Partitioning** window, select the mount point.
2. Click the **Device Type** drop-down menu and select **LVM**. The **Volume Group** drop-down menu is displayed with the newly-created volume group name.



NOTE

The configuration dialog does not allow you to specify the size of the volume group's physical extents. The size is always be set to the default value of 4 MiB. If you want to create a volume group with different physical extents, create it manually by switching to an interactive shell and using the **vgcreate** command, or use a Kickstart file with the **volgroup --pesize=size** command.

6.8.7. Configuring an LVM logical volume

This section describes how to configure a newly-created LVM logical volume.

Procedure

1. From the left-hand pane of the **Manual Partitioning** window, select the mount point.
2. Click the **Device Type** drop-down menu and select **LVM**. The **Volume Group** drop-down menu is displayed with the newly-created volume group name.
3. Click **Modify** to configure the newly-created volume group.

The **Configure Volume Group** dialog box opens.



NOTE

The configuration dialog does not allow you to specify the size of the volume group's physical extents. The size is always set to the default value of 4 MiB. If you want to create a volume group with different physical extents, create it manually by switching to an interactive shell and using the **vgcreate** command, or use a Kickstart file with the **volgroup --pesize=size** command.

4. From the **RAID Level** drop-down menu, select the required RAID level.
The available RAID levels are the same as with actual RAID devices.
5. Select the **Encrypt** check box to mark the volume group for encryption.
6. From the **Size policy** drop-down menu, select the size policy for the volume group.
The available policy options are:
 - **Automatic:** The size of the volume group is set automatically so that it is large enough to contain the configured logical volumes. This is optimal if you do not need free space within the volume group.
 - **As large as possible:** The volume group is created with maximum size, regardless of the size of the configured logical volumes it contains. This is optimal if you plan to keep most of your data on LVM and later need to increase the size of some existing logical volumes, or if you need to create additional logical volumes within this group.
 - **Fixed:** You can set an exact size of the volume group. Any configured logical volumes must then fit within this fixed size. This is useful if you know exactly how large you need the volume group to be.
7. Click **Save** to apply the settings and return to the **Manual Partitioning** window.
8. Click **Update Settings** to save your changes
9. Click **Done** to return to the **Installation Summary** window.



WARNING

Placing the **/boot** partition on an LVM volume is not supported.

6.9. STARTING THE INSTALLATION PROGRAM

Starting the installation program requires the configuration of your root password and user settings.

6.9.1. Beginning installation


When you have started the installation process, it is not possible to go back to the **Installation Summary** window and change any settings. To change settings, you must wait for the installation process to finish, reboot your system, log in, and change your settings on the installed system.

Prerequisites

- You have completed all configuration steps in [Section 6.3, “Installation summary”](#).
- The **Installation Summary** window is open.

Procedure

1. From the **Installation Summary** window, click **Begin Installation**. The **Configuration** window opens and the installation process starts.
Two user setting options, **Root Password** (mandatory) and **User Creation** (optional) are available.




IMPORTANT

Before you finish the installation and reboot, either remove the media (CD, DVD, or a USB drive) used to start the installation, or verify that your system tries to boot from the hard drive before attempting removable media. Otherwise, your system starts the installation program again, instead of the installed system.

6.9.2. Configuring a root password

Configuring a **root** password is required to finish the installation process and to log into the administrator (also known as superuser or root) account that is used for system administration tasks. These tasks include installing and updating software packages and changing system-wide configuration such as network and firewall settings, storage options, and adding or modifying users, groups and file permissions.



IMPORTANT

- Use one or both of the following ways to gain root privileges to the installed system:
 - Use a root account
 - Create a user account with administrative privileges (member of the wheel group). The **root** account is always created during the installation. Only switch to the administrator account when you need to perform a task that requires administrator access.



WARNING

The **root** account has complete control over the system. If unauthorized personnel gain access to the account, they can access or delete users' personal files.

Procedure

1. From the **Configuration** window, click **Root Password**. The **Root Password** window opens.
2. Type your password in the **Root Password** field. For security purposes, the characters are displayed as dots.
 - a. The requirements and recommendations for creating a strong root password are:
 - i. *Must* be at least eight characters long
 - ii. May contain numbers, letters (upper and lower case) and symbols
 - iii. Is case-sensitive
3. Type the same password in the **Confirm** field.
4. Click **Done** to confirm your root password and return to [Section 6.9.1, “Beginning installation”](#).



NOTE

If you proceeded with a weak password, you must click **Done** twice.

6.9.3. Creating a user account

It is recommended that you create a user account to finish the installation. If you do not create a user account, you must log in to the system as **root** directly, which is **not** recommended.

Procedure


1. From the **Configuration** window, click **User Creation**. The **Create User** window opens.
2. Type the user account name into the **Full name** field, for example: John Smith.
3. Type the username into the **User name** field, for example: jsmith.



NOTE

The **User name** is used to log in from a command line; if you install a graphical environment, then your graphical login manager uses the **Full name**.

4. Select the **Make this user administrator** check box if the user requires administrative rights (it is added into the **wheel** group).



IMPORTANT

An administrator user can use the **sudo** command to perform tasks that are only available to **root** using the user password, instead of the **root** password. This may be more convenient, but it can also cause a security risk.

5. Select the **Require a password to use this account** check box.

**WARNING**

If you give administrator privileges to a user, verify that the account is password protected. Never give a user administrator privileges without assigning a password to the account.

6. Type a password into the **Password** field.
7. Type the same password into the **Confirm password** field.
8. Click **Save Changes** to apply the changes and return to the **Configuration** window.
9. Click **Reboot** to reboot and log in to your Red Hat Enterprise Linux 8 system.

6.9.3.1. Configuring Advanced User Settings

You can change the default settings for the user account in the **Advanced User Configuration** dialog box.

Procedure

1. Change the details in the **Home directory** field, if required. The field is populated by default with `/home/username`.
2. In the **User and Groups IDs** section you can:
 - a. Select the **Specify a user ID manually** check box and use the **+** or **-** to enter the required value.

**NOTE**

The default value is 1000. UIDs 0-999 are reserved by the system so they can not be assigned to a user.

- b. Select the **Specify a group ID manually** check box and use the **+** or **-** to enter the required value.

**NOTE**

The default group name is the same as the user name, and its default GID is 1000. GIDs 0-999 are reserved by the system so they can not be assigned to a user's group.

3. Specify additional groups as a comma-separated list in the **Group Membership** field. Groups that do not already exist are created; you can specify custom GIDs for them in parentheses. If you do not specify a custom GID for a new group, it is assigned automatically.

**NOTE**

The user account created always has one default group membership (the user's default group with an ID set in the **Specify a group ID manually** field).

4. Click **Save Changes** to apply the updates and return to the **Configuration** window.

6.9.4. Installation complete

Congratulations! Your Red Hat Enterprise Linux 8 installation is complete! Remove any installation media if it is not ejected automatically upon reboot.

After your system's normal power-up sequence completes, Red Hat Enterprise Linux 8 loads and starts. When complete, a GUI login window (or if the X Window System is not installed, a **login:** prompt) is displayed.

If your system was installed with the X Window System, applications to set up your system are launched the first time you start your Red Hat Enterprise Linux 8 system. These applications guide you through initial configuration and you can set your system time and date, register your machine with Red Hat Network, and more.

6.9.5. UEFI Secure Boot for RHEL 8

UEFI Secure Boot requires that the operating system kernel is signed with a recognized private key. For Red Hat Enterprise Linux 8, the kernel is signed with a Red Hat Beta-specific private key, which is different from the standard Red Hat key used in a General Availability release.

Red Hat Enterprise Linux 8 cannot boot if your hardware does not recognize the Beta private key. To use UEFI Secure Boot with the Beta release, add the Red Hat Beta public key to your system using the Machine Owner Key (MOK) facility.

6.9.6. Adding a custom private key for UEFI Secure Boot

This section provides instructions on how to add a private key for UEFI Secure Boot on Red Hat Enterprise Linux 8.

Prerequisite

Disable UEFI Secure Boot on the system and install Red Hat Enterprise Linux 8.

Procedure

1. Enroll the Red Hat CA public key in the system's Machine Owner Key (MOK) list:

```
# kr=$(uname -r)
# mokutil --import /usr/share/doc/kernel-keys/${kr%}/kernel-signing-ca.cer
```

`/${kr%}` is replaced by the string without the platform architecture suffix - for example, **3.10.0-686.el7**.

2. Enter a password.
3. Reboot the system.

4. During system startup, you are prompted to confirm the key request. Select **Yes** and enter the password.
The system reboots and the key is imported into the system's firmware.
5. Enable Secure Boot on the system.

**WARNING**

Remove the imported Beta public key if you install a General Availability (GA) release of Red Hat Enterprise Linux 8 or if you install a different operating system.

CHAPTER 7. POST-INSTALLATION TASKS

Post-installation tasks covers finalizing the installation of Red Hat Enterprise Linux 8.

- Completing initial setup
- Registering the system
- Securing your system

7.1. COMPLETING INITIAL SETUP

This section provides information on how to complete the initial setup in **graphical mode** on Red Hat Enterprise Linux 8. The Initial Setup window will open for the first time you boot your system after installation completes. Please, notice that this information may change depending on what has already been previously configured. At a minimum, the **Licensing** and **Subscription Manager** options are displayed.

Prerequisite

- You have completed the graphical installation as per the recommended workflow described on [Section 6.1, “Graphical installation workflow”](#).
- You have an active Red Hat Subscription prepared for the registration process.

Procedure

1. From the **Initial Setup** window, select **Licensing**.
The **License Agreement** screen opens and displays the licensing terms for Red Hat Enterprise Linux.
2. Review the license agreement and select the **I accept the license agreement** checkbox.



NOTE

The license agreement must be accepted. Exiting **Initial Setup** without completing this step causes the system to restart. Once the system completes the restart, you are prompted to accept the agreement again.

3. Click **Done** to apply the settings and return to the **Initial Setup** window.



NOTE

If you did not configure the network settings previously, you will not be able to register your system immediately. In this case, click the **Finish Configuration** button to finish the configuration. Red Hat Enterprise Linux 8 will start and you will be able to login, activate access to the network and register your system then. [Section 7.3, “Registering your system using the Subscription Manager User Interface”](#) If you have activated access to the internet previously, as described at [Section 6.6.3, “Network and host name”](#), you will be able to register your system immediately. For that, follow the instructions below:

4. From the **Initial Setup** window, select **Subscription Manager**.

5. The **Subscription Manager** graphical interface opens and displays the option you are going to register: `"subscription.rhsm.redhat.com"` and click the **Next** button.
 6. Insert Login and Password and click the **Register** button.
 7. Confirm the Subscription details and click the **Attach** button.
 8. You must receive the following message confirmation: **Registration with Red Hat Subscription Management is Done!**
 9. Click on button **Done** and you will be redirected to the **Initial Setup** screen.
 10. Click on button **Finish Configuration**. The system will lead you to the login screen.
 11. Configure your system. See *Configuring basic system settings* for more information.
- Additional resources**

There are four methods to register your system:

- During installation using Initial Setup.
- After installation using the command line. See [Section 7.2, “Registering your system using the command line”](#) for more information.
- After installation using the Subscription Manager user interface. See [Section 7.3, “Registering your system using the Subscription Manager User Interface”](#).
- After installation using Registration Assistant. Registration Assistant is designed to help you choose the most suitable registration option for your Red Hat Enterprise Linux environment. See <https://access.redhat.com/labs/registrationassistant/> for more information .

7.2. REGISTERING YOUR SYSTEM USING THE COMMAND LINE

This section provides instructions on how to register your Red Hat Enterprise Linux 8 system using the command line.



NOTE

When auto-attaching a system, the subscription service checks if the system is physical or virtual, as well as how many sockets are on the system. A physical system usually consumes two entitlements, a virtual system usually consumes one. One entitlement is consumed per two sockets on a system.

Prerequisites

- You have an active, non-evaluation Red Hat Enterprise Linux subscription.
- Your Red Hat subscription status is verified.
- You have not previously received a Red Hat Enterprise Linux 8 subscription.
- You have activated your subscription before attempting to download entitlements from the Customer Portal. You need an entitlement for each instance you plan to use. Red Hat Customer Service is available if you need help activating your subscription.
- You have successfully installed a Red Hat Enterprise Linux 8 system and logged into it.

Procedure

1. Open a terminal window and type the following command register a subscription:

```
# subscription-manager register
```

2. Enter your Customer Portal Credentials:

```
# Registering to: subscription.rhsm.redhat.com:443/subscription
# Username: USERNAME
# Password: PASSWORD
```

3. When the subscription is successfully registered, you should see an output similar to the following example:

```
# The system has been registered with ID: 123456abcdef
# The registered system name is: localhost.localdomain
```

4. Attach the system to an entitlement that matches the host system's architecture:

```
# subscription-manager attach
```

5. When the subscription is successfully attached, you should see an output similar to the following example:

```
Installed Product Current Status:
Product Name: Red Hat Enterprise Linux for 86_64 High Touch Beta
Status: Subscribed
```



NOTE

You can also register Red Hat Enterprise Linux 8 in a graphical interface by logging into the system with the **root** user and using the Subscription Manager GUI.

7.3. REGISTERING YOUR SYSTEM USING THE SUBSCRIPTION MANAGER USER INTERFACE

This section provides instructions on how to register your Red Hat Enterprise Linux 8 system using the Subscription Manager User Interface and thus be able to receive updates and access to package repositories.

Prerequisites

- You have completed the graphical installation as per the recommended workflow described on [Section 6.1, “Graphical installation workflow”](#).
- You have an active Red Hat Subscription.

Procedure

1. Login into your system.

2. Click on the network connection icon on the system menu on the right side of the top bar to connect to your network.
3. On the top left side of the screen, click on **Activities** to show **All Application** icon.
4. Find and Run the application **Red Hat Subscription Manager**.
5. You are prompted to insert your administrator password into Authentication Required window.

**NOTE**

Authentication is required to perform privileged tasks into the system.

6. A window showing Subscription Status, Status Purpose and Installed Product will be prompted.
7. Click on **Register** button.
8. Enter your Customer Portal Credentials and click on button **Register**.

As a result, the system shows the Subscription Information Status as Registered.

7.4. REGISTRATION ASSISTANT

Registration Assistant is designed to help you choose the most suitable registration option for your Red Hat Enterprise Linux environment.

See <https://access.redhat.com/labs/registrationassistant/> for more information

7.5. SECURING YOUR SYSTEM

The following steps are the security-related procedures that should be performed immediately after installation of Red Hat Enterprise Linux.

Prerequisites

- You have completed the graphical installation as per the recommended workflow described on [Section 6.1, “Graphical installation workflow”](#).

Procedure

1. Update your system. enter the following command as root:

```
# dnf update
```

2. Even though the firewall service, `firewalld`, is automatically enabled with the installation of Red Hat Enterprise Linux, there are scenarios where it might be explicitly disabled, for example in the kickstart configuration. In such a case, it is recommended to consider re-enabling the firewall.

To start `firewalld` enter the following commands as root:

```
# systemctl start firewalld
# systemctl enable firewalld
```

3. To enhance security, disable services you do not need. For example, if there are no printers installed on your computer, disable the cups service using the following command:

```
# systemctl disable cups
```

To review active services, enter the following command:

```
$ systemctl list-units | grep service
```

Additional resources

- For further information see XXX

APPENDIX A. SYSTEM REQUIREMENTS REFERENCE

This section provides information and guidelines for hardware, installation target, system, memory, and RAID requirements when installing Red Hat Enterprise Linux.

A.1. CHECK HARDWARE COMPATIBILITY

Red Hat works closely with hardware vendors on supported hardware.

- To verify that your hardware is supported, see the Red Hat Hardware Compatibility List, available at <https://access.redhat.com/ecosystem/search/#/category/Server>.
- To view supported memory sizes or CPU counts, see <https://access.redhat.com/articles/rhel-limits> for information.

A.2. REVIEW SUPPORTED INSTALLATION TARGETS

An installation target is a storage device that stores Red Hat Enterprise Linux and boots the system. Red Hat Enterprise Linux supports the following installation targets for AMD, Intel, and ARM systems:

- Storage connected by a standard internal interface, such as SCSI, SATA, or SAS
- BIOS/firmware RAID devices
- NVDIMM devices in sector mode on the Intel64 and AMD64 architectures, supported by the `nd_pmem` driver.
- Fibre Channel Host Bus Adapters and multipath devices. Some can require vendor-provided drivers.
- Xen block devices on Intel processors in Xen virtual machines.
- VirtIO block devices on Intel processors in KVM virtual machines.

Red Hat does not support installation to USB drives or SD memory cards. For information on support for third-party virtualization technologies, see the [Red Hat Hardware Compatibility List](#).

A.3. RECORD SYSTEM SPECIFICATIONS

The Red Hat Enterprise Linux installation program automatically detects and installs your system's hardware, so you should not have to supply any specific system information. However, for certain Red Hat Enterprise Linux installation scenarios, it is recommended that you record system specifications for future reference. These scenarios include:

Installing RHEL with a customized partition layout

Record: The model numbers, sizes, types, and interfaces of the hard drives attached to the system. For example, Seagate ST3320613AS 320 GB on SATA0, Western Digital WD7500AAKS 750 GB on SATA1.

Installing RHEL as an additional operating system on an existing system

Record: Partitions used on the system. This information can include file system types, device node names, file system labels, and sizes, and allows you to identify specific partitions during the partitioning process. If one of the operating systems is a Unix operating system, Red Hat Enterprise Linux may report the device names differently. Additional information can be found by executing the equivalent of the `mount` command and the `blkid` command, and in the `/etc/fstab` file.

If multiple operating systems are installed, the Red Hat Enterprise Linux installation program attempts to automatically detect them, and to configure boot loader to boot them. You can manually configure additional operating systems if they are not detected automatically. See *Configuring boot loader* in [Section 6.5, “Software settings”](#) for more information.

Installing RHEL from an image on a local hard drive

Record: The hard drive and directory that holds the image.

Installing RHEL from a network location

If the network has to be configured manually, that is, DHCP is not used.

Record: * IP address * Netmask * Gateway IP address * Server IP addresses, if required

Contact your network administrator if you need assistance with networking requirements.

Installing RHEL on an iSCSI target

Record: The location of the iSCSI target. Depending on your network, you may need a CHAP user name and password, and a reverse CHAP user name and password.

Installing RHEL if the system is part of a domain

Verify that the domain name is supplied by the DHCP server. If it is not, enter the domain name during installation.

A.4. CHECK DISK AND MEMORY REQUIREMENTS

If several operating systems are installed, it is important that you verify that the allocated disk space is separate from the disk space required by Red Hat Enterprise Linux.



NOTE

For AMD, Intel 64-bit, and 64-bit ARM, at least two partitions (**/** and **swap**) must be dedicated to Red Hat Enterprise Linux.

You must have a minimum of 10 GiB of available disk space. See [Appendix B, Partitioning reference](#) for more information.

Table A.1. Minimum RAM requirements

Installation type	Recommended minimum RAM
Local media installation (USB, DVD)	768 MiB
NFS network installation	768 MiB
HTTP, HTTPS, or FTP network installation	1.5 GiB

**NOTE**

It is possible to complete the installation with less memory than the recommended minimum requirements. The exact requirements depend on your environment and installation path. It is recommended that you test various configurations to determine the minimum required RAM for your environment. Installing Red Hat Enterprise Linux using a Kickstart file has the same recommended minimum RAM requirements as a standard installation. However, additional RAM may be required if your Kickstart file includes commands that require additional memory, or write data to the RAM disk. See the *Installing_RHEL_as_an_experienced_user* guide for more information.

A.5. REVIEW RAID REQUIREMENTS

It is important to understand how storage technologies are configured and how support for them may have changed between major versions of Red Hat Enterprise Linux.

Hardware RAID

Any RAID functions provided by the mainboard of your computer, or attached controller cards, need to be configured before you begin the installation process. Each active RAID array appears as one drive within Red Hat Enterprise Linux.

Software RAID

On systems with more than one hard drive, you can use the Red Hat Enterprise Linux installation program to operate several of the drives as a Linux software RAID array. With a software RAID array, RAID functions are controlled by the operating system rather than the dedicated hardware.

**NOTE**

When a pre-existing RAID array's member devices are all unpartitioned disks/drives, the installation program treats the array as a disk and there is no method to remove the array.

USB Disks

You can connect and configure external USB storage after installation. Most devices are recognized by the kernel, but some devices may not be recognized. If it is not a requirement to configure these disks during installation, disconnect them to avoid potential problems.

NVDIMM devices

To use a Non-Volatile Dual In-line Memory Module (NVDIMM) device as storage, the following conditions must be satisfied:

- Version of Red Hat Enterprise Linux is 7.6 or later.
- The architecture of the system is Intel 64 or AMD64.
- The device is configured to sector mode. Anaconda can reconfigure NVDIMM devices to this mode.
- The device must be supported by the `nd_pmem` driver.

Bootting from a NVDIMM device is possible under the following additional conditions:

- The system uses UEFI.

- The device must be supported by firmware available on the system, or by a UEFI driver. The UEFI driver may be loaded from an option ROM of the device itself.
- The device must be made available under a namespace.

To take advantage of the high performance of NVDIMM devices during booting, place the **/boot** and **/boot/efi** directories on the device. See **LINK TO BE ADDED** for more information.



NOTE

The Execute-in-place (XIP) feature of NVDIMM devices is not supported during booting and the kernel is loaded into conventional memory.

Considerations for Intel BIOS RAID Sets

Red Hat Enterprise Linux uses **mdraid** for installing on Intel BIOS RAID sets. These sets are automatically detected during the boot process and their device node paths can change across several booting processes. For this reason, local modifications to the **/etc/fstab**, **/etc/crypttab** or other configuration files that refer to the devices by their device node paths may not work in Red Hat Enterprise Linux. It is recommended that you replace device node paths (such as **/dev/sda**) with file system labels or device UUIDs. You can find the file system labels and device UUIDs using the **blkid** command.

APPENDIX B. PARTITIONING REFERENCE

B.1. SUPPORTED DEVICE TYPES

Standard partition

A standard partition can contain a file system or swap space. Standard partitions are most commonly used for **/boot** and the **BIOS Boot** and **EFI System partitions**. LVM logical volumes are recommended for most other uses.

LVM

Choosing **LVM** (or Logical Volume Management) as the Device Type creates an LVM logical volume. If no LVM volume group currently exists, one is automatically created to contain the new volume, if one already exists, the volume is assigned to it. LVM can improve performance when using physical disks and allows for advanced setups such as using multiple physical disks for one mount point, and setting up software RAID for increased performance, reliability, or both.

LVM Thin Provisioning

Using thin provisioning, you can manage a storage pool of free space, known as a thin pool, which can be allocated to an arbitrary number of devices when needed by applications. You can dynamically expand the pool when needed for cost-effective allocation of storage space.

B.2. SUPPORTED FILE SYSTEMS

This section describes the file systems available in the installer.

xfs

XFS is a highly scalable, high-performance file system that supports file systems up to 16 exabytes (approximately 16 million terabytes), files up to 8 exabytes (approximately 8 million terabytes), and directory structures containing tens of millions of entries. XFS also supports metadata journaling, which facilitates quicker crash recovery. The maximum supported size of a single XFS file system is 500 TB. XFS is the default and recommended file system on Red Hat Enterprise Linux.

ext4

The ext4 file system is based on the ext3 file system and features a number of improvements. These include support for larger file systems and larger files, faster and more efficient allocation of disk space, no limit on the number of subdirectories within a directory, faster file system checking, and more robust journaling. The maximum supported size of a single ext4 file system is 50 TB.

ext3

The ext3 file system is based on the ext2 file system and has one main advantage - journaling. Using a journaling file system reduces time spent recovering a file system after a crash, as there is no need to check the file system for metadata consistency by running the fsck utility every time a crash occurs.

ext2

An ext2 file system supports standard Unix file types, including regular files, directories, or symbolic links. It provides the ability to assign long file names, up to 255 characters.

swap

Swap partitions are used to support virtual memory. In other words, data is written to a swap partition when there is not enough RAM to store the data your system is processing.

vfat

The VFAT file system is a Linux file system that is compatible with Microsoft Windows long file names on the FAT file system.

BIOS Boot

A very small partition required for booting from a device with a GUID partition table (GPT) on BIOS systems and UEFI systems in BIOS compatibility mode.

EFI System Partition

A small partition required for booting a device with a GUID partition table (GPT) on a UEFI system.

B.3. SUPPORTED RAID TYPES

RAID stands for Redundant Array of Independent Disks, a technology which allows you to combine multiple physical disks into logical units. Some setups are designed to enhance performance at the cost of reliability, while others will improve reliability at the cost of requiring more disks for the same amount of available space.

This section describes supported software RAID types which you can use with LVM and LVM Thin Provisioning to set up the installed system's storage.

None

No RAID array will be set up.

RAID0

Performance - Distributes data across multiple disks. Level 0 RAID offers increased performance over standard partitions and can be used to pool the storage of multiple disks into one large virtual device. Note that Level 0 RAID offers no redundancy and that the failure of one device in the array destroys data in the entire array. RAID 0 requires at least two disks.

RAID1

Redundancy - Mirrors all data from one partition onto one or more other disks. Additional devices in the array provide increasing levels of redundancy. RAID 1 requires at least two disks.

RAID4

Error checking - Distributes data across multiple disks and uses one disk in the array to store parity information which safeguards the array in case any disk within the array fails. Because all parity information is stored on one disk, access to this disk creates a "bottleneck" in the array's performance. Level 4 RAID requires at least three disks.

RAID5

Distributed error checking - Distributes data and parity information across multiple disks. Level 5 RAID therefore offers the performance advantages of distributing data across multiple disks, but does not share the performance bottleneck of level 4 RAID because the parity information is also distributed through the array. RAID 5 requires at least three disks.

RAID6

Redundant error checking - Level 6 RAID is similar to level 5 RAID, but instead of storing only one set of parity data, they store two sets. RAID 6 requires at least four disks.

RAID10

Performance and redundancy - Level 10 RAID is a nested RAID or hybrid RAID. They are constructed by distributing data over mirrored sets of disks. For example, a level 10 RAID array constructed from four RAID partitions consists of two mirrored pairs of striped partitions. RAID 10 requires at least four disks.

B.4. RECOMMENDED PARTITIONING SCHEME

Red Hat recommends that you create separate file systems at the following mount points:

- **/boot**

- `/` (root)
- `/home`
- `swap`
- `/boot/efi`

`/boot` partition - recommended size at least 1 GiB

The partition mounted on `/boot` contains the operating system kernel, which allows your system to boot Red Hat Enterprise Linux 8, along with files used during the bootstrap process. Due to the limitations of most firmwares, creating a small partition to hold these is recommended. In most scenarios, a 1 GiB boot partition is adequate. Unlike other mount points, using an LVM volume for `/boot` is not possible - `/boot` must be located on a separate disk partition.



WARNING

Normally, the `/boot` partition is created automatically by the installation program. However, if the `/` (root) partition is larger than 2 TiB and (U)EFI is used for booting, you need to create a separate `/boot` partition that is smaller than 2 TiB to boot the machine successfully.



NOTE

If you have a RAID card, be aware that some BIOS types do not support booting from the RAID card. In such a case, the `/boot` partition must be created on a partition outside of the RAID array, such as on a separate hard drive.

root - recommended size of 10 GiB

This is where `/`, or the root directory, is located. The root directory is the top-level of the directory structure. By default, all files are written to this file system unless a different file system is mounted in the path being written to, for example, `/boot` or `/home`.

While a 5 GiB root file system allows you to install a minimal installation, it is recommended to allocate at least 10 GiB so that you can install as many package groups as you want.



IMPORTANT

Do not confuse the `/` directory with the `/root` directory. The `/root` directory is the home directory of the root user. The `/root` directory is sometimes referred to as *slash root* to distinguish it from the root directory.

`/home` - recommended size at least 1 GiB

To store user data separately from system data, create a dedicated file system for the `/home` directory. Base the file system size on the amount of data that is stored locally,

number of users, and so on. You can upgrade or reinstall Red Hat Enterprise Linux 8 without erasing user data files. If you select automatic partitioning, it is recommended to have at least 55 GiB of disk space available for the installation, to ensure that the `/home` file system is created.

swap partition - recommended size at least 1 GB

Swap file systems support virtual memory; data is written to a swap file system when there is not enough RAM to store the data your system is processing. Swap size is a function of system memory workload, not total system memory and therefore is not equal to the total system memory size. It is important to analyze what applications a system will be running and the load those applications will serve in order to determine the system memory workload. Application providers and developers can provide guidance.

When the system runs out of swap space, the kernel terminates processes as the system RAM memory is exhausted. Configuring too much swap space results in storage devices being allocated but idle and is a poor use of resources. Too much swap space can also hide memory leaks. The maximum size for a swap partition and other additional information can be found in the **mkswap(8)** manual page.

The following table provides the recommended size of a swap partition depending on the amount of RAM in your system and if you want sufficient memory for your system to hibernate. If you let the installation program partition your system automatically, the swap partition size is established using these guidelines. Automatic partitioning setup assumes hibernation is not in use. The maximum size of the swap partition is limited to 10 percent of the total size of the hard drive, and the installation program cannot create swap partitions more than 128 GB in size. To set up enough swap space to allow for hibernation, or if you want to set the swap partition size to more than 10 percent of the system's storage space, or more than 128 GB, you must edit the partitioning layout manually.

/boot/efi partition - recommended size of 200 MiB

UEFI-based AMD, Intel, and ARM require a 200 MiB EFI system partition. The recommended minimum size is 200 MiB, the default size is 600 MiB, and the maximum size is 600 MiB. BIOS systems do not require an EFI system partition.

Table B.1. Recommended System Swap Space

Amount of RAM in the system	Recommended swap space	Recommended swap space if allowing for hibernation
Less than 2 GB	2 times the amount of RAM	3 times the amount of RAM
2 GB - 8 GB	Equal to the amount of RAM	2 times the amount of RAM
8 GB - 64 GB	4 GB to 0.5 times the amount of RAM	1.5 times the amount of RAM
More than 64 GB	Workload dependent (at least 4GB)	Hibernation not recommended

At the border between each range, for example, a system with 2 GB, 8 GB, or 64 GB of system RAM, discretion can be exercised with regard to chosen swap space and hibernation support. If your system resources allow for it, increasing the swap space can lead to better performance.

Distributing swap space over multiple storage devices - particularly on systems with fast drives, controllers and interfaces - also improves swap space performance.

Many systems have more partitions and volumes than the minimum required. Choose partitions based on your particular system needs.



NOTE

Only assign storage capacity to those partitions you require immediately. You can allocate free space at any time, to meet needs as they occur.



NOTE

If you are unsure about how to configure partitions, accept the automatic default partition layout provided by the installation program.

B.5. ADVICE ON PARTITIONS

There is no best way to partition every system; the optimal setup depends on how you plan to use the system being installed. However, the following tips may help you find the optimal layout for your needs:

- Create partitions that have specific requirements first, for example, if a particular partition must be on a specific disk.
- Consider encrypting any partitions and volumes which might contain sensitive data. Encryption prevents unauthorized people from accessing the data on the partitions, even if they have access to the physical storage device. In most cases, you should at least encrypt the **/home** partition, which contains user data.
- In some cases, creating separate mount points for directories other than **/**, **/boot** and **/home** may be useful; for example, on a server running a **MySQL** database, having a separate mount point for **/var/lib/mysql** will allow you to preserve the database during a reinstallation without having to restore it from backup afterwards. However, having unnecessary separate mount points will make storage administration more difficult.
- Some special restrictions apply to certain directories with regards on which partitioning layouts can they be placed. Notably, the **/boot** directory must always be on a physical partition (not on an LVM volume).
- If you are new to Linux, consider reviewing the *Linux Filesystem Hierarchy Standard* at http://refspecs.linuxfoundation.org/FHS_2.3/fhs-2.3.html for information about various system directories and their contents.
- Each kernel installed on your system requires approximately 20 MB on the **/boot** partition. The default partition size of 1 GB for **/boot** should suffice for most common uses; increase the size of this partition if you plan to keep many kernels installed at the same time.
- The **/var** directory holds content for a number of applications, including the **Apache** web server, and is used by the **DNF** package manager to temporarily store downloaded package updates. Make sure that the partition or volume containing **/var** has at least 3 GB.
- The contents of the **/var** directory usually change very often. This may cause problems with older solid state drives (SSDs), as they can handle a lower number of read/write cycles before becoming unusable. If your system root is on an SSD, consider creating a separate mount point for **/var** on a classic (platter) HDD.

- The **/usr** directory holds the majority of software on a typical Red Hat Enterprise Linux installation. The partition or volume containing this directory should therefore be at least 5 GB for minimal installations, and at least 10 GB for installations with a graphical environment.
- If **/usr** or **/var** is partitioned separately from the rest of the root volume, the boot process becomes much more complex because these directories contain boot-critical components. In some situations, such as when these directories are placed on an iSCSI drive or an FCoE location, the system may either be unable to boot, or it may hang with a **Device is busy** error when powering off or rebooting.
This limitation only applies to **/usr** or **/var**, not to directories below them. For example, a separate partition for **/var/www** will work without issues.
- Consider leaving a portion of the space in an LVM volume group unallocated. This unallocated space gives you flexibility if your space requirements change but you do not wish to remove data from other volumes. You can also select the **LVM Thin Provisioning** device type for the partition to have the unused space handled automatically by the volume.
- The size of an XFS file system can not be reduced - if you need to make a partition or volume with this file system smaller, you must back up your data, destroy the file system, and create a new, smaller one in its place. Therefore, if you expect needing to manipulate your partitioning layout later, you should use the ext4 file system instead.
- Use Logical Volume Management (LVM) if you anticipate expanding your storage by adding more hard drives after the installation. With LVM, you can create physical volumes on the new drives, and then assign them to any volume group and logical volume as you see fit - for example, you can easily expand your system's **/home** (or any other directory residing on a logical volume).
- Creating a BIOS Boot partition or an EFI System Partition may be necessary, depending on your system's firmware, boot drive size, and boot drive disk label. See [Section B.4, "Recommended partitioning scheme"](#) for information about these partitions. Note that the graphical installer will not let you create a BIOS Boot or EFI System Partition if your system does **not** require one - in that case, they will be hidden from the menu.
- If you need to make any changes to your storage configuration after the installation, Red Hat Enterprise Linux repositories offer several different tools which can help you do this. If you prefer a command line tool, try **system-storage-manager**.

APPENDIX C. TROUBLESHOOTING

The following sections cover various troubleshooting information which may be helpful when diagnosing installation issues.

C.1. CONSOLES AND LOGGING DURING INSTALLATION

The Red Hat Enterprise Linux installer uses the **tmux** terminal multiplexer to display and control several windows you can use in addition to the main interface. Each of these windows serves a different purpose - they display several different logs, which can be used to troubleshoot any issues during the installation, and one of the windows provides an interactive shell prompt with **root** privileges, unless this prompt was specifically disabled using a boot option or a Kickstart command.



NOTE

In general, there is no reason to leave the default graphical installation environment unless you need to diagnose an installation problem.

The terminal multiplexer is running in virtual console 1. To switch from the actual installation environment to **tmux**, press **Ctrl+Alt+F1**. To go back to the main installation interface which runs in virtual console 6, press **Ctrl+Alt+F6**.



NOTE

If you choose text mode installation, you will start in virtual console 1 (**tmux**), and switching to console 6 will open a shell prompt instead of a graphical interface.

The console running **tmux** has 5 available windows; their contents are described in the table below, along with keyboard shortcuts used to access them. Note that the keyboard shortcuts are two-part: first press **Ctrl+b**, then release both keys, and press the number key for the window you want to use.

You can also use **Ctrl+b n** and **Ctrl+b p** to switch to the next or previous **tmux** window, respectively.

Table C.1. Available tmux windows

Shortcut	Contents
Ctrl+b 1	Main installation program window. Contains text-based prompts (during text mode installation or if you use VNC direct mode), and also some debugging information.
Ctrl+b 2	Interactive shell prompt with root privileges.
Ctrl+b 3	Installation log; displays messages stored in /tmp/anaconda.log .
Ctrl+b 4	Storage log; displays messages related storage devices from kernel and system services, stored in /tmp/storage.log .

Shortcut	Contents
Ctrl+b 5	Program log; displays messages from other system utilities, stored in /tmp/program.log .

C.2. SAVING SCREENSHOTS

You can press **Shift+Print Screen** at any time during the graphical installation to capture the current screen. These screenshots are saved to **/tmp/anaconda-screenshots**.

C.3. RESUMING AN INTERRUPTED DOWNLOAD ATTEMPT

You can resume an interrupted download using the **curl** command.

Procedure

1. Refresh the download page in the Customer Portal; log in again if necessary.
2. Copy the new download link.
3. Download the ISO image from the new link. Add the **--continue-at -** option to automatically resume the download:

```
#curl --output directory-path/filename.iso
'new_copied_link_location' --continue-at -
```

4. Use a checksum utility such as **sha256sum** to verify the integrity of the image file after the download finishes:

```
$ sha256sum rhel-server-8.0-x86_64-dvd.iso
`85a...46c rhel-server-8.0-x86_64-dvd.iso`
```

Compare the output with reference checksums provided on the Red Hat Enterprise Linux **Download** web page.

Example C.1. Resuming an interrupted download attempt

The following is an example of a **curl** command for a partially downloaded ISO image:

```
$ curl --output _rhel-server-8.0-x86_64-dvd.iso
'https://access.cdn.redhat.com//content/origin/files/sha256/85/85a...46c
/rhel-server-8.0-x86_64-dvd.iso?_auth=141...963' --continue-at -
```

PART II. INSTALLING RED HAT ENTERPRISE LINUX ON IBM POWER

This section describes how to install Red Hat Enterprise Linux on the IBM Power architecture.

CHAPTER 8. PREPARATION FOR IBM POWER SYSTEMS SERVERS

The following text describes how to install Red Hat Enterprise Linux on the IBM Power architecture.

8.1. PREREQUISITES

Include modules here.

8.2. INSTALLATION OR UPGRADING

8.2.1. Planning for Installation on IBM Power Systems

While automated in-place upgrades are now supported, the support is currently limited to AMD64 and Intel 64 systems. Therefore, any existing installation of a previous release of Red Hat Enterprise Linux on an IBM Power Systems server, a clean install must be performed to migrate to Red Hat Enterprise Linux 8. A clean install is performed by

- Purchasing, or ensuring an active subscription
- Backing up all data from the system
- Formatting disk partitions
- Performing an installation of Red Hat Enterprise Linux from installation media
- Restoring any user data

Additional Resources

- For more information on checking your Red Hat Enterprise Linux subscription, see <add link to checking subscription>

8.3. PREPARATION FOR IBM POWER SYSTEMS SERVERS

Non-partitioned systems do not need any pre-installation setup. For systems using the HVSI serial console, hook up the console to the T2 serial port.

The steps to create the partition and start the installation on partitioned systems are largely the same. Using the HMC create partition wizard,

- Create the partition at the HMC
- Assign CPU and memory resources
- Assign SCSI and Ethernet resources—either virtual or native.

Additional Resources

- For more information on creating the partition, see the Partitioning for Linux with an HMC in the IBM Systems Hardware Information Center at:
http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/topic/iphbi_p5/iphbibook.pdf

8.4. IBM INSTALLATION TOOLS

8.4.1. IBM Installation Toolkit

[IBM Installation Toolkit](#) is an optional utility that speeds up the installation of Linux on IBM Power Systems and is especially helpful for those unfamiliar with Linux. You can use the [IBM Installation Toolkit](#) to:

- Install and configure Linux on a non-virtualized IBM Power Systems server.
- Install and configure Linux on servers with previously-configured logical partitions (LPARs, also known as virtualized servers).
- Install IBM service and productivity tools on a new or previously installed Linux system. The IBM service and productivity tools include dynamic logical partition (DLPAR) utilities.
- Upgrade system firmware level on IBM Power Systems servers.
- Perform diagnostics or maintenance operations on previously installed systems.
- Migrate a LAMP server (software stack) and application data from a System x to a System p system. A LAMP server is a bundle of open source software. LAMP is an acronym for Linux, **A**pache **H**TTP **S**erver, **M**ySQL relational database, and the PHP (or sometimes Perl, or Python) language.

8.5. POWERLINUX PRODUCTIVITY TOOLS

PowerLinux service and productivity tools is an optional set of tools that includes hardware service diagnostic aids, productivity tools, and installation aids for Linux operating systems on IBM servers based on POWER7, POWER6, POWER5, and POWER4 technology.

Additional resources


- Documentation for the service and productivity tools is available in the [Linux Information Center](#)

8.6. SUPPORTED INSTALLATION TARGETS

An installation target is a storage device that stores Red Hat Enterprise Linux and boots the system. Red Hat Enterprise Linux supports the following installation targets for AMD64 and Intel 64 systems:

- Storage connected by a standard internal interface, such as SCSI, SATA, or SAS.
- Fibre Channel Host Bus Adapters and multipath devices. Some of these devices may require vendor-provided drivers.
- Virtualized installation on IBM Power Systems servers when using Virtual SCSI (vSCSI) adapters in virtual client LPARs.

Red Hat does not support installation to USB drives or SD memory cards. For support information on third-party virtualization technologies, see the *Red Hat Hardware Compatibility List*, available online at <https://hardware.redhat.com>.



IMPORTANT

On IBM Power Systems servers, the eHEA module fails to initialize if 16GB *huge pages* are assigned to a system or partition and the kernel command line does not contain the huge page parameters. When you perform a network installation through an IBM eHEA ethernet adapter, you cannot assign huge pages to the system or perform a partition during the installation. Use *large pages* instead.

Additional resources

- For information about the support for third-party virtualization technologies, see the *Red Hat Hardware Compatibility List*, available online at <https://hardware.redhat.com>.

8.7. SYSTEM SPECIFICATIONS LIST

The installation program automatically detects and installs software compatible with the computer's hardware, and normally it is not necessary to supply details about your system. However, when performing certain types of installation, it is important to know specific details about your hardware. For this reason, it is recommended that you record the following system specifications for reference during the installation, depending on your installation type.

- If you plan to use a customized partition layout, record:
 - The model numbers, sizes, types, and interfaces of the hard drives attached to the system. For example, Seagate ST3320613AS 320 GB on SATA0, Western Digital WD7500AAKS 750 GB on SATA1. This will allow you to identify specific hard drives during the partitioning process.
- If you are installing Red Hat Enterprise Linux as an additional operating system on an existing system, record:
 - Information about the partitions used on the system. This information can include file system types, device node names, file system labels, and sizes. This will allow you to identify specific partitions during the partitioning process. Remember that different operating systems identify partitions and drives differently, therefore even if the other operating system is a Unix operating system, the device names can be reported by Red Hat Enterprise Linux differently. This information can usually be found by executing the equivalent of the **mount** command and **blkid** command and in the **/etc/fstab** file. If other operating systems are already installed, the Red Hat Enterprise Linux 8 installation program attempts to automatically detect and configure to boot them. You can manually configure any additional operating systems if they are not detected properly. For more information, see
- If you plan to install from an image on a local hard drive:
 - The hard drive and directory that holds the image.
- If you plan to install from a network location:
 - The make and model numbers of the network adapters on your system. For example, Netgear GA311. This will allow you to identify adapters when manually configuring the network.
 - IP, DHCP, and BOOTP addresses
 - Netmask

- Gateway IP address
- One or more name server IP addresses (DNS)
- The location of the installation source on an FTP server, HTTP (web) server, HTTPS (web) server, or NFS server.
If any of these networking requirements or terms are unfamiliar to you, contact your network administrator for assistance.
- If you plan to install on an iSCSI target:
 - The location of the iSCSI target. Depending on your network, you might also need a CHAP user name and password, and perhaps a reverse CHAP user name and password.
- If your computer is part of a domain:
 - You should verify that the domain name will be supplied by the DHCP server. If not, you will need to input the domain name manually during installation.

8.8. DISK SPACE AND MEMORY REQUIREMENTS

Red Hat Enterprise Linux requires minimum the following amount of RAM:

Installation type	Minimum required RAM
Local media installation (USB, DVD)	1,280 MiB
NFS network installation	1,280 MiB
HTTP, HTTPS, or FTP network installation	1,664 MiB



NOTE

Installing Red Hat Enterprise Linux using a Kickstart file has the same minimum RAM requirements as a manual installation. However, if you use a Kickstart file that runs commands which require additional memory or write data to the RAM disk, additional RAM might be necessary.

Red Hat Enterprise Linux requires minimum the following amount of disk space:

Amount	Requirements
10 GB of space in either unpartitioned disk space or in partitions which can be deleted.	The disk space used by Red Hat Enterprise Linux must be separate from the disk space used by other operating systems that might be installed on your system.

For IBM Power Systems servers, at least three partitions must be dedicated to Red Hat Enterprise Linux.

Partition	Type
/	Root
swap	Swap
PReP	Boot

Additional Resources

- For more information about disk partitions, see
- For more information about the minimum requirements and technology limits of Red Hat Enterprise Linux 8, see the [Red Hat Enterprise Linux technology capabilities and limits](#) article on the Red Hat Customer Portal.
- For more information on partition and disk space recommendations, see the recommended partitioning sizes discussed in

8.9. RAID AND OTHER DISK DEVICES

Some storage technologies require special consideration when using Red Hat Enterprise Linux. Generally, it is important to understand how these technologies are configured, visible to Red Hat Enterprise Linux, and how support for them might have changed between major versions.

8.9.1. RAID and Other Disk Devices

- Redundant Array of Independent Disks (RAID) allows a group, or array, of drives to act as a single device.

8.9.1.1. Hardware RAID

- Any RAID functions provided by the mainboard of your computer, or attached controller cards, need to be configured before you begin the installation process. Each active RAID array appears as one drive within Red Hat Enterprise Linux.

8.9.1.2. Software RAID

- On systems with more than one hard drive, you can use the Red Hat Enterprise Linux installation program to operate several of the drives as a Linux software RAID array. With a software RAID array, RAID functions are controlled by the operating system rather than dedicated hardware.



NOTE

When a pre-existing RAID array's member devices are all unpartitioned disks/drives, the installer will treat the array itself as a disk and will not provide a way to remove the array.

8.9.1.3. USB Disks

- You can connect and configure external USB storage after installation. Most such devices are recognized by the kernel and available for use at that time.

Some USB drives might not be recognized by the installation program. If configuration of these disks at installation time is not vital, disconnect them to avoid potential problems.

Additional Resources

- For more information on software RAID functions, see

8.10. CHOOSE AN INSTALLATION BOOT METHOD

There are several methods to boot the Red Hat Enterprise Linux installation program. The method you choose depends upon your installation media.



NOTE

Installation media must remain mounted throughout the installation process, including during the execution of the `%post` section of a kickstart file.

Full installation DVD or USB drive

A full installation DVD or USB drive is created using the Binary DVD ISO image. It can be used as both a boot device and as an installation source for installing software packages.

Minimal installation CD, DVD or USB flash drive

A minimal installation CD, DVD, or USB flash drive is created using the boot ISO image, which only contains the minimum files necessary to boot the system and start the installation program. This boot option requires an installation source that contains the required software packages.

PXE Server

A *preboot execution environment* (PXE) server allows the installation program to boot over the network. After you boot the system, you complete the installation from a different installation source, such as a local hard drive or a location on a network.

Additional Resources

- For instructions on how to create a full installation DVD or USB drive, see [Section 4.6, “Create installation media”](#) for more information.
- For instructions on how to create a boot CDs, DVDs and USB flash drive, see [Section 4.7, “Prepare an installation source”](#) for more information.
- For more information on PXE servers, see <XX TO BE ADDED>

8.11. AUTOMATING THE INSTALLATION WITH KICKSTART

{ProductNameNAME} 8 offers a way to partially or fully automate the installation process using a *Kickstart file*. Kickstart files contain answers to all questions normally asked by the installation program, such as what time zone do you want the system to use, how should the drives be partitioned or which packages should be installed. Providing a prepared Kickstart file at the beginning of the installation therefore allows you to perform the entire installation (or parts of it) automatically, without need for any intervention from the user. This is especially useful when deploying Red Hat Enterprise Linux on a large number of systems at once.

In addition to allowing you to automate the installation, Kickstart files also provide more options

regarding software selection. When installing Red Hat Enterprise Linux manually using the graphical installation interface, your software selection is limited to pre-defined environments and add-ons. A Kickstart file allows you to install or remove individual packages as well.

Additional Resources

- For instructions about creating a Kickstart file and using it to automate the installation, see

8.12. RELATED INFORMATION

- A bulleted list of links to other material closely related to the contents of the concept module.
- For more details on writing assemblies, see the [Modular Documentation Reference Guide](#).
- Use a consistent system for file names, IDs, and titles. For tips, see *Anchor Names and File Names* in [Modular Documentation Reference Guide](#).

PART III. INSTALLING RED HAT ENTERPRISE LINUX ON IBM Z

This section describes how to install Red Hat Enterprise Linux on the IBM Z architecture.

CHAPTER 9. PREPARING FOR INSTALLATION ON IBM Z

This paragraph is the assembly introduction. It explains what the user will accomplish by working through the modules in the assembly and sets the context for the user story the assembly is based on. Can include more than one paragraph. Consider using the information from the user story.

9.1. PREREQUISITES

- A bulleted list of conditions that must be satisfied before the user starts following this assembly.
- You can also link to other modules or assemblies the user must follow before starting this assembly.
- Delete the section title and bullets if the assembly has no prerequisites.

9.2. OVERVIEW OF THE IBM Z INSTALLATION PROCESS

You can install Red Hat Enterprise Linux on IBM Z interactively or in unattended mode. Installation on IBM Z differs from installation on other architectures in that it is typically performed over a network and not from local media. The installation consists of two phases:

1. Booting the installation

- Connect with the mainframe
- Perform an initial program load (IPL), or boot, from the medium containing the installation program.

2. Anaconda

Use the **Anaconda** installation program to:

- Configure the network
- Specify language support
- Specify installation source
- Specify software packages to be installed
- Perform the rest of the installation

Additional resources

- A bulleted list of links to other material closely related to the contents of the concept module.
- For more details on writing concept modules, see the [Modular Documentation Reference Guide](#).
- Use a consistent system for file names, IDs, and titles. For tips, see *Anchor Names and File Names* in [Modular Documentation Reference Guide](#).

9.3. PLANNING FOR INSTALLATION ON IBM Z

9.3.1. Pre-installation

Red Hat Enterprise Linux 8 runs on z 13 or later IBM mainframe systems.

The installation process assumes that you are familiar with the IBM Z and can set up *logical partitions* (LPARs) and z/VM guest virtual machines.

For installation of Red Hat Enterprise Linux on IBM Z, Red Hat supports Direct Access Storage Device (DASD) and Fiber Channel Protocol (FCP) storage devices.

Pre-installation decisions

- Whether the operating system is to be run on an LPAR or as a z/VM guest operating system.
- If swap space is needed, and how much. Although it is recommended to assign enough memory to a z/VM guest virtual machine and let z/VM do the necessary swapping, there are cases where the amount of required RAM is hard to predict. Such instances should be examined on a case-by-case basis.
- Network configuration. Red Hat Enterprise Linux 8 for IBM Z supports the following network devices:
 - Real and virtual *Open Systems Adapter*(OSA)
 - Real and virtual HiperSockets
 - *LAN channel station*(LCS) for real OSA

Disk space

You will need to calculate and allocate sufficient disk space on DASDs or SCSI disks.

- A minimum of 10 GB is needed for a server installation, 20 GB if you want to install all packages.
- Disk space is also required for any application data. After the installation, you can add or delete more DASD or SCSI disk partitions.
- The disk space used by the newly installed Red Hat Enterprise Linux system (the Linux instance) must be separate from the disk space used by other operating systems you have installed on your system.

RAM

You will have to ensure enough RAM is available.

- 1 GB is recommended for the Linux instance. With some tuning, an instance might run with as little as 512 MB RAM.
- If installing from nfs, 1 GB is sufficient. However, if installing from an http/ftp source, 1.5 GB is needed.
- Running at 512 MB in text mode can be done only when installing from nfs.



NOTE

When initializing swap space on a Fixed Block Architecture (FBA) DASD using the **SWAPGEN** utility, the **FBAPART** option must be used.

Additional Resources

- For additional information on IBM Z, see <http://www.ibm.com/systems/z>.

9.4. INSTALLING UNDER Z/VM

Use the **x3270** or **c3270** terminal emulator, to log in to z/VM from other Linux systems, or use the IBM 3270 terminal emulator on the IBM Z Hardware Management Console (HMC). If you are running Microsoft Windows operating system, there are several options available, and can be found through an internet search. A free native Windows port of **c3270** called **wc3270** also exists.

When installing under z/VM, you can boot from:

- The z/VM virtual reader
 - A DASD or an FCP-attached SCSI device prepared with the **zipl** boot loader
 - An FCP-attached SCSI DVD drive
1. Log on to the z/VM guest virtual machine chosen for the Linux installation.



NOTE

If your 3270 connection is interrupted and you cannot log in again because the previous session is still active, you can replace the old session with a new one by entering the following command on the z/VM logon screen:

```
logon user here
```

Replace *user* with the name of the z/VM guest virtual machine. Depending on whether an external security manager, for example RACF, is used, the logon command might vary.

If you are not already running **CMS** (single-user operating system shipped with z/VM) in your guest, boot it now by entering the command:

```
cp ipl cms
```

Be sure not to use CMS disks such as your A disk (often device number 0191) as installation targets. To find out which disks are in use by CMS, use the following query:

```
query disk
```

You can use the following CP (z/VM Control Program, which is the z/VM hypervisor) query commands to find out about the device configuration of your z/VM guest virtual machine:

- Query the available main memory, which is called *storage* in IBM Z terminology. Your guest should have at least 1 GB of main memory.

```
cp query virtual storage
```

- Query available network devices by type:

```
osa
```

OSA - CHPID type OSD, real or virtual (VSWITCH or GuestLAN), both in QDIO mode

```
hsi
```


HiperSockets - CHPID type IQD, real or virtual (GuestLAN type Hipers)

lcs

LCS - CHPID type OSE

For example, to query all of the network device types mentioned above, run:

```
cp query virtual osa
```

- Query available DASDs. Only those that are flagged **RW** for read-write mode can be used as installation targets:

```
cp query virtual dasd
```

- Query available FCP channels:

```
cp query virtual fcp
```

9.5. USING PARAMETER AND CONFIGURATION FILES ON IBM Z

The IBM Z architecture can use a customized parameter file to pass boot parameters to the kernel and the installation program.

You need to change the parameter file if you want to:

- Install unattended with Kickstart.
- Choose non-default installation settings that are not accessible through the installation program's interactive user interface, such as rescue mode.

The parameter file can be used to set up networking non-interactively before the installation program (**Anaconda**) starts.

The kernel parameter file is limited to 895 characters plus an end-of-line character. The parameter file can be variable or fixed record format. Fixed record format increases the file size by padding each line up to the record length. Should you encounter problems with the installation program not recognizing all specified parameters in LPAR environments, you can try to put all parameters in one single line or start and end each line with a space character.


The parameter file contains kernel parameters, such as **ro**, and parameters for the installation process, such as **vncpassword=test** or **vnc**.

9.6. REQUIRED CONFIGURATION FILE PARAMETERS ON IBM Z

Several parameters are required and must be included in the parameter file. These parameters are also provided in the file **generic.prm** in directory **images/** of the installation DVD.

- **ro**
Mounts the root file system, which is a RAM disk, read-only.
- **ramdisk_size=size**
Modifies the memory size reserved for the RAM disk to ensure that the Red Hat Enterprise Linux installation program fits within it. For example: **ramdisk_size=40000**.

The **generic.prm** file also contains the additional parameter **cio_ignore=all,!condev**. This setting speeds up boot and device detection on systems with many devices. The installation program transparently handles the activation of ignored devices.



IMPORTANT

To avoid installation problems arising from **cio_ignore** support not being implemented throughout the entire stack, adapt the **cio_ignore=** parameter value to your system or remove the parameter entirely from your parameter file used for booting (IPL) the installation program.

9.7. IBM ZVM CONFIGURATION FILE

Under z/VM, you can use a configuration file on a CMS-formatted disk. The purpose of the CMS configuration file is to save space in the parameter file by moving the parameters that configure the initial network setup, the DASD, and the FCP specification out of the parameter file.

Each line of the CMS configuration file contains a single variable and its associated value, in the following shell-style syntax: **variable=value**.

You must also add the **CMSDASD** and **CMSCONFFILE** parameters to the parameter file. These parameters point the installation program to the configuration file:

CMSDASD=cmsdasd_address

Where *cmsdasd_address* is the device number of a CMS-formatted disk that contains the configuration file. This is usually the CMS user's **A** disk.

For example: **CMSDASD=191**

CMSCONFFILE=configuration_file

Where *configuration_file* is the name of the configuration file. This value must be specified in lower case. It is specified in a Linux file name format: **CMS_file_name.CMS_file_type**.

The CMS file **REDHAT CONF** is specified as **redhat.conf**. The CMS file name and the file type can each be from one to eight characters that follow the CMS conventions.

For example: **CMSCONFFILE=redhat.conf**

9.8. INSTALLATION NETWORK PARAMETERS ON IBM Z

These parameters can be used to automatically set up the preliminary network, and can be defined in the CMS configuration file. These parameters are the only parameters that can also be used in a CMS configuration file. All other parameters in other sections must be specified in the parameter file.

NETTYPE="type"

Where *type* must be one of the following: **qeth**, **lcs**, or **ctc**. The default is **qeth**.

Choose **lcs** for:

- OSA-2 Ethernet/Token Ring
- OSA-Express Fast Ethernet in non-QDIO mode
- OSA-Express High Speed Token Ring in non-QDIO mode

- Gigabit Ethernet in non-QDIO mode
Choose **qeth** for:
- OSA-Express Fast Ethernet
- Gigabit Ethernet (including 1000Base-T)
- High Speed Token Ring
- HiperSockets
- ATM (running Ethernet LAN emulation)

SUBCHANNELS="device_bus_IDs"

Where *device_bus_IDs* is a comma-separated list of two or three device bus IDs. The IDs must be specified in lowercase.

Provides required device bus IDs for the various network interfaces:

```
qeth:
SUBCHANNELS="read_device_bus_id,write_device_bus_id,data_device_bus_id"
lcs or ctc: SUBCHANNELS="read_device_bus_id,write_device_bus_id"
```

For example (a sample qeth SUBCHANNEL statement):

```
SUBCHANNELS="0.0.f5f0,0.0.f5f1,0.0.f5f2"
```

PORTNAME="osa_portname" PORTNAME="lcs_portnumber"

This variable supports OSA devices operating in qdio mode or in non-qdio mode.

When using qdio mode (**NETTYPE="qeth"**), *osa_portname* is the portname specified on the OSA device when operating in qeth mode.

When using non-qdio mode (**NETTYPE="lcs"**), *lcs_portnumber* is used to pass the relative port number as a decimal integer in the range of 0 through 15.

PORTNO="portnumber"

You can add either **PORTNO="0"** (to use port 0) or **PORTNO="1"** (to use port 1 of OSA features with two ports per CHPID) to the CMS configuration file to avoid being prompted for the mode.

LAYER2="value"

Where *value* can be **0** or **1**.

Use **LAYER2="0"** to operate an OSA or HiperSockets device in layer 3 mode (**NETTYPE="qeth"**).

Use **LAYER2="1"** for layer 2 mode. For virtual network devices under z/VM this setting must match the definition of the GuestLAN or VSWITCH to which the device is coupled.

To use network services that operate on layer 2 (the Data Link Layer or its MAC sublayer) such as DHCP, layer 2 mode is a good choice.

The qeth device driver default for OSA devices is now layer 2 mode. To continue using the previous default of layer 3 mode, set **LAYER2="0"** explicitly.

VSWITCH="value"

Where *value* can be **0** or **1**.

Specify **VSWITCH="1"** when connecting to a z/VM VSWITCH or GuestLAN, or **VSWITCH="0"** (or nothing at all) when using directly attached real OSA or directly attached real HiperSockets.

MACADDR="MAC_address"

If you specify **LAYER2="1"** and **VSWITCH="0"**, you can optionally use this parameter to specify a MAC address. Linux requires six colon-separated octets as pairs lower case hex digits - for example, **MACADDR=62:a3:18:e7:bc:5f**. Note that this is different from the notation used by z/VM.

If you specify **LAYER2="1"** and **VSWITCH="1"**, you must not specify the **MACADDR**, because z/VM assigns a unique MAC address to virtual network devices in layer 2 mode.

CTCPROT="value"

Where *value* can be **0**, **1**, or **3**.

Specifies the CTC protocol for **NETTYPE="ctc"**. The default is **0**.

HOSTNAME="string"

Where *string* is the host name of the newly-installed Linux instance.

IPADDR="IP"

Where *IP* is the IP address of the new Linux instance.

NETMASK="netmask"

Where *netmask* is the netmask.

The netmask supports the syntax of a prefix integer (from 1 to 32) as specified in IPv4 *classless interdomain routing* (CIDR). For example, you can specify **24** instead of **255.255.255.0**, or **20** instead of **255.255.240.0**.

GATEWAY="gw"

Where *gw* is the gateway IP address for this network device.

MTU="mtu"

Where *mtu* is the *Maximum Transmission Unit* (MTU) for this network device.

DNS="server1:server2:additional_server_terms:serverN"

Where "server1:server2:additional_server_terms:serverN" is a list of DNS servers, separated by colons. For example:

```
DNS="10.1.2.3:10.3.2.1"
```

SEARCHDNS="domain1:domain2:additional_dns_terms:domainN"

Where "domain1:domain2:additional_dns_terms:domainN" is a list of the search domains, separated by colons. For example:

```
SEARCHDNS="subdomain.domain:domain"
```

You only need to specify **SEARCHDNS=** if you specify the **DNS=** parameter.

DASD=

Defines the DASD or range of DASDs to configure for the installation.

The installation program supports a comma-separated list of device bus IDs, or ranges of device bus IDs with the optional attributes **ro**, **diag**, **erplog**, and **failfast**. Optionally, you can abbreviate device bus IDs to device numbers with leading zeros stripped. Any optional attributes should be

separated by colons and enclosed in parentheses. Optional attributes follow a device bus ID or a range of device bus IDs.

The only supported global option is **autodetect**. This does not support the specification of non-existent DASDs to reserve kernel device names for later addition of DASDs. Use persistent DASD device names (for example **/dev/disk/by-path/...**) to enable transparent addition of disks later. Other global options such as **probeonly**, **nopav**, or **nofcx** are not supported by the installation program.

Only specify those DASDs that need to be installed on your system. All unformatted DASDs specified here must be formatted after a confirmation later on in the installation program.

Add any data DASDs that are not needed for the root file system or the **/boot** partition after installation.

For example:

```
DASD="eb1c,0.0.a000-0.0.a003,eb10-eb14(diag),0.0.ab1c(ro:diag)"
```

For FCP-only environments, remove the **DASD=** option from the CMS configuration file to indicate no DASD is present.

FCP_ *n*="device_bus_ID WWPN FCP_LUN"

Where:

- *n* is typically an integer value (for example **FCP_1** or **FCP_2**) but could be any string with alphabetic or numeric characters or underscores.
- *device_bus_ID* specifies the device bus ID of the FCP device representing the *host bus adapter* (HBA) (for example **0.0.fc00** for device fc00).
- *WWPN* is the world wide port name used for routing (often in conjunction with multipathing) and is as a 16-digit hex value (for example **0x50050763050b073d**).
- *FCP_LUN* refers to the storage logical unit identifier and is specified as a 16-digit hexadecimal value padded with zeroes to the right (for example **0x4020400100000000**). These variables can be used on systems with FCP devices to activate FCP LUNs such as SCSI disks. Additional FCP LUNs can be activated during the installation interactively or by means of a Kickstart file. An example value looks similar to the following:

```
FCP_1="0.0.fc00 0x50050763050b073d 0x4020400100000000"
```



IMPORTANT

Each of the values used in the FCP parameters (for example **FCP_1** or **FCP_2**) are site-specific and are normally supplied by the FCP storage administrator.

The installation program prompts you for any required parameters not specified in the parameter or configuration file except for FCP_*n*.

9.9. PARAMETERS FOR KICKSTART INSTALLATIONS ON IBM Z

The following parameters can be defined in a parameter file but do not work in a CMS configuration file.

inst.ks=URL

References a Kickstart file, which usually resides on the network for Linux installations on IBM Z. Replace *URL* with the full path including the file name of the Kickstart file. This parameter activates automatic installation with Kickstart.

RUNKS=value



IMPORTANT

This parameter is deprecated. If you use it in a Kickstart file, it will be ignored. Only the **inst.ks** parameter is necessary to start a Kickstart installation on IBM Z.

Where *value* is defined as *1* if you want to run the loader automatically on the Linux console without having to log in over the network with SSH. To use **RUNKS=1**, the console must either support full-screen or the **inst.cmdline** option (below) should be used. The latter applies for the 3270 terminal under z/VM or the operating system messages console for LPAR. We recommend **RUNKS=1** for fully automatic installations with Kickstart. When **RUNKS=1** is set, the installation program automatically continues in case of parameter errors and does not interrupt unattended installations by prompting for user interaction.

Leave out the parameter or specify **RUNKS=0** otherwise.

inst.cmdline

When this option is specified, output on line-mode terminals (such as 3270 under z/VM or operating system messages for LPAR) becomes readable, as the installation program disables escape terminal sequences that are only applicable to UNIX-like consoles. This requires installation with a Kickstart file that answers all questions, because the installation program does not support interactive user input in cmdline mode.

Ensure that your Kickstart file contains all required parameters before you use the **inst.cmdline** option. If a required command is missing, the installation will fail.

9.10. MISCELLANEOUS PARAMETERS ON IBM Z

The following parameters can be defined in a parameter file but do not work in a CMS configuration file.

rd.live.check

Turns on testing of an ISO-based installation source; for example, when booted from an FCP-attached DVD or using **inst.repo=** with an ISO on local hard disk or mounted with NFS.

nopath

Disables support for multipath devices.

proxy=[protocol://][username[:password]@]host[:port]

Specify a proxy to use with installation over HTTP, HTTPS, or FTP.

inst.rescue

Boot into a rescue system running from a RAM disk that can be used to fix and restore an installed system.

inst.stage2=URL

Specifies a path to an **install.img** file instead of to an installation source. Otherwise, follows the

same syntax as **inst.repo=**. If **inst.stage2** is specified, it typically takes precedence over other methods of finding **install.img**. However, if **Anaconda** finds **install.img** on local media, the **inst.stage2** URL will be ignored.

If **inst.stage2** is not specified and **install.img** cannot be found locally, **Anaconda** looks to the location given by **inst.repo=** or **method=**.

If only **inst.stage2=** is given without **inst.repo=** or **method=**, **Anaconda** uses whatever repos the installed system would have enabled by default for installation.

Use the option multiple times to specify multiple HTTP, HTTPS or FTP sources. The HTTP, HTTPS or FTP paths are then tried sequentially until one succeeds:

```
inst.stage2=host1/install.img
inst.stage2=host2/install.img
inst.stage3=host3/install.img
```

inst.syslog=IP/hostname[:port]

Sends log messages to a remote syslog server.

The boot parameters described here are the most useful for installations and trouble shooting on IBM Z, but only a subset of those that influence the installation program.

9.11. SAMPLE PARAMETER FILE AND CMS CONFIGURATION FILE ON IBM Z

To change the parameter file, begin by extending the shipped **generic.prm** file.

Example of **generic.prm** file:

```
ro ramdisk_size=40000 cio_ignore=all,!condev
CMSDASD="191" CMSCONFFILE="redhat.conf"
vnc
inst.repo=http://example.com/path/to/repository
```

Example of **redhat.conf** file configuring a QETH network device (pointed to by **CMSCONFFILE** in **generic.prm**):

```
NETTYPE="qeth"
SUBCHANNELS="0.0.0600,0.0.0601,0.0.0602"
PORTNAME="FOOBAR"
PORTNO="0"
LAYER2="1"
MACADDR="02:00:be:3a:01:f3"
HOSTNAME="foobar.systemz.example.com"
IPADDR="192.168.17.115"
NETMASK="255.255.255.0"
GATEWAY="192.168.17.254"
DNS="192.168.17.1"
SEARCHDNS="systemz.example.com:example.com"
DASD="200-203"
```

CHAPTER 10. CONFIGURING A LINUX INSTANCE ON IBM Z

This paragraph is the assembly introduction. It explains what the user will accomplish by working through the modules in the assembly and sets the context for the user story the assembly is based on. Can include more than one paragraph. Consider using the information from the user story.

10.1. PREREQUISITES

- A bulleted list of conditions that must be satisfied before the user starts following this assembly.
- You can also link to other modules or assemblies the user must follow before starting this assembly.
- Delete the section title and bullets if the assembly has no prerequisites.

10.2. DASDs THAT ARE PART OF THE ROOT FILE SYSTEM

The file you have to modify to add DASDs that are part of the root file system has changed in Red Hat Enterprise Linux 8. Instead of editing the `/etc/zip1.conf` file, the new file to be edited, and its location, may be found by running the following commands:

```
# machine_id=$(cat /etc/machine-id)
# kernel_version=$(uname -r)
# ls /boot/loader/entries/$machine_id-$kernel_version.conf
```

There is one boot option to activate DASDs early in the boot process: **rd.dasd=**. This option takes a comma-separated list as input. The list contains a device bus ID and optional additional parameters consisting of key-value pairs that correspond to DASD **sysfs** attributes.

Below is an example of the `/boot/loader/entries/4ab74e52867b4f998e73e06cf23fd761-4.18.0-32.el8.s390x.conf` file for a system that uses physical volumes on partitions of two DASDs for an LVM volume group **vg_devel1** that contains a logical volume **lv_root** for the root file system.

```
title Red Hat Enterprise Linux (4.18.0-32.el8.s390x) 8.0 (Ootpa)
version 4.18.0-32.el8.s390x
linux /boot/vmlinuz-4.18.0-32.el8.s390x
initrd /boot/initramfs-4.18.0-32.el8.s390x.img
options root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.lvm.lv=vg_devel1/lv_root
rd.lvm.lv=vg_devel1/lv_swap cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0
id rhel-20181027190514-4.18.0-32.el8.s390x
grub_users $grub_users
grub_arg --unrestricted
grub_class kernel
```

To add another physical volume on a partition of a third DASD with device bus ID **0.0.202b**. To do this, add **rd.dasd=0.0.202b** to the parameters line of your boot kernel in **zip1.conf**:

```
title Red Hat Enterprise Linux (4.18.0-32.el8.s390x) 8.0 (Ootpa)
version 4.18.0-32.el8.s390x
linux /boot/vmlinuz-4.18.0-32.el8.s390x
```



```

initrd /boot/initramfs-4.18.0-32.el8.s390x.img
options root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.dasd=0.0.202b
rd.lvm.lv=vg_devel1/lv_root rd.lvm.lv=vg_devel1/lv_swap
cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0
id rhel-20181027190514-4.18.0-32.el8.s390x
grub_users $grub_users
grub_arg --unrestricted
grub_class kernel

```



WARNING

Make sure the length of the kernel command line in the configuration file does not exceed 896 bytes. Otherwise, the boot loader cannot be saved, and the installation fails.

Run **zipl** to apply the changes of the configuration file for the next IPL:

```

# zipl -V
Using config file '/etc/zipl.conf'
Using BLS config file
'/boot/loader/entries/4ab74e52867b4f998e73e06cf23fd761-4.18.0-
32.el8.s390x.conf'
Target device information
  Device.....: 5e:00
  Partition.....: 5e:01
  Device name.....: dasda
  Device driver name.....: dasd
  DASD device number.....: 0201
  Type.....: disk partition
  Disk layout.....: ECKD/compatible disk layout
  Geometry - heads.....: 15
  Geometry - sectors.....: 12
  Geometry - cylinders.....: 13356
  Geometry - start.....: 24
  File system block size.....: 4096
  Physical block size.....: 4096
  Device size in physical blocks...: 262152
Building bootmap in '/boot'
Building menu 'zipl-automatic-menu'
Adding #1: IPL section '4.18.0-32.el8.s390x' (default)
  initial ramdisk...: /boot/initramfs-4.18.0-32.el8.s390x.img
  kernel image.....: /boot/vmlinuz-4.18.0-32.el8.s390x
  kernel parmline...: 'root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.dasd=0.0.202b
rd.lvm.lv=vg_devel1/lv_root rd.lvm.lv=vg_devel1/lv_swap
cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0'
  component address:

```

```

kernel image....: 0x00010000-0x0049afff
parmline.....: 0x0049b000-0x0049bfff
initial ramdisk.: 0x004a0000-0x01a26fff
internal loader.: 0x0000a000-0x0000cfff
Preparing boot menu
Interactive prompt.....: enabled
Menu timeout.....: 5 seconds
Default configuration...: '4.18.0-32.el8.s390x'
Preparing boot device: dasda (0201).
Syncing disks...
Done.

```

10.3. FCP LUNS THAT ARE PART OF THE ROOT FILE SYSTEM

The only file you have to modify for adding FCP LUNs that are part of the root file system has changed in Red Hat Enterprise Linux 8. Instead of editing the `/etc/zip1.conf` file, the new file to be edited, and its location, may be found by running the following commands:

```

# machine_id=$(cat /etc/machine-id)
# kernel_version=$(uname -r)
# ls /boot/loader/entries/$machine_id-$kernel_version.conf

```

Red Hat Enterprise Linux provides a parameter to activate FCP LUNs early in the boot process: **rd.zfcp=**. The value is a comma-separated list containing the device bus ID, the WWPN as 16 digit hexadecimal number prefixed with **0x**, and the FCP LUN prefixed with **0x** and padded with zeroes to the right to have 16 hexadecimal digits.

Below is an example of the `/boot/loader/entries/4ab74e52867b4f998e73e06cf23fd761-4.18.0-32.el8.s390x.conf` file for a system that uses physical volumes on partitions of two FCP LUNs for an LVM volume group **vg_devel1** that contains a logical volume **lv_root** for the root file system. For simplicity, the example shows a configuration without multipathing.

```

title Red Hat Enterprise Linux (4.18.0-32.el8.s390x) 8.0 (Ootpa)
version 4.18.0-32.el8.s390x
linux /boot/vmlinuz-4.18.0-32.el8.s390x
initrd /boot/initramfs-4.18.0-32.el8.s390x.img
options root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.lvm.lv=vg_devel1/lv_root
rd.lvm.lv=vg_devel1/lv_swap cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0
id rhel-20181027190514-4.18.0-32.el8.s390x
grub_users $grub_users
grub_arg --unrestricted
grub_class kernel

```

To add another physical volume on a partition of a third FCP LUN with device bus ID **0.0.fc00**, WWPN **0x5105074308c212e9** and FCP LUN **0x401040a300000000**, add **rd.zfcp=0.0.fc00,0x5105074308c212e9,0x401040a300000000** to the parameters line of your boot kernel in `zip1.conf`. For example:

```

title Red Hat Enterprise Linux (4.18.0-32.el8.s390x) 8.0 (Ootpa)
version 4.18.0-32.el8.s390x
linux /boot/vmlinuz-4.18.0-32.el8.s390x
initrd /boot/initramfs-4.18.0-32.el8.s390x.img

```

```
options root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.dasd=0.0.202b
rd.lvm.lv=vg_devel1/lv_root rd.lvm.lv=vg_devel1/lv_swap
cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0
id rhel-20181027190514-4.18.0-32.el8.s390x
grub_users $grub_users
grub_arg --unrestricted
grub_class kernel
```

**WARNING**

Make sure the length of the kernel command line in the configuration file does not exceed 896 bytes. Otherwise, the boot loader cannot be saved, and the installation fails.

Run **zipl** to apply the changes of the configuration file for the next IPL:

```
# zipl -V
Using config file '/etc/zipl.conf'
Using BLS config file
'/boot/loader/entries/4ab74e52867b4f998e73e06cf23fd761-4.18.0-
32.el8.s390x.conf'
Target device information
  Device.....: 5e:00
  Partition.....: 5e:01
  Device name.....: dasda
  Device driver name.....: dasd
  DASD device number.....: 0201
  Type.....: disk partition
  Disk layout.....: ECKD/compatible disk layout
  Geometry - heads.....: 15
  Geometry - sectors.....: 12
  Geometry - cylinders.....: 13356
  Geometry - start.....: 24
  File system block size.....: 4096
  Physical block size.....: 4096
  Device size in physical blocks...: 262152
Building bootmap in '/boot'
Building menu 'zipl-automatic-menu'
Adding #1: IPL section '4.18.0-32.el8.s390x' (default)
  initial ramdisk...: /boot/initramfs-4.18.0-32.el8.s390x.img
  kernel image.....: /boot/vmlinuz-4.18.0-32.el8.s390x
  kernel parmline...: 'root=/dev/mapper/vg_devel1-lv_root crashkernel=auto
rd.dasd=0.0.0200 rd.dasd=0.0.0207 rd.dasd=0.0.202b
rd.lvm.lv=vg_devel1/lv_root rd.lvm.lv=vg_devel1/lv_swap
cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0'
  component address:
    kernel image....: 0x00010000-0x0049afff
```

```

    parmline.....: 0x0049b000-0x0049bfff
    initial ramdisk.: 0x004a0000-0x01a26fff
    internal loader.: 0x0000a000-0x0000cfff
Preparing boot menu
Interactive prompt.....: enabled
Menu timeout.....: 5 seconds
Default configuration...: '4.18.0-32.el8.s390x'
Preparing boot device: dasda (0201).
Syncing disks...
Done.

```

10.4. ADDING A QETH DEVICE

The **qeth** network device driver supports IBM Z OSA-Express features in QDIO mode, HiperSockets, z/VM guest LAN, and z/VM VSWITCH.

The **qeth** device driver assigns the same interface name for Ethernet and Hipersockets devices: **enclbus_ID**. The bus ID is composed of the channel subsystem ID, subchannel set ID, and device number, and does not contain leading zeros and dots. For example **enca00** for a device with the bus ID **0.0.0a00**.

10.5. DYNAMICALLY ADDING A QETH DEVICE

To add a **qeth** device dynamically, follow these steps:

1. Determine whether the **qeth** device driver modules are loaded. The following example shows loaded **qeth** modules:

```

# lsmod | grep qeth
qeth_l3          69632  0
qeth_l2          49152  1
qeth             131072  2 qeth_l3,qeth_l2
qdio             65536  3 qeth,qeth_l3,qeth_l2
ccwgroup         20480  1 qeth

```

If the output of the **lsmod** command shows that the **qeth** modules are not loaded, run the **modprobe** command to load them:

```
# modprobe qeth
```

2. Use the **cio_ignore** utility to remove the network channels from the list of ignored devices and make them visible to Linux:

```
# cio_ignore -r
read_device_bus_id,write_device_bus_id,data_device_bus_id
```

Replace *read_device_bus_id*, *write_device_bus_id*, *data_device_bus_id* with the three device bus IDs representing a network device. For example, if the *read_device_bus_id* is **0.0.f500**, the *write_device_bus_id* is **0.0.f501**, and the *data_device_bus_id* is **0.0.f502**:

```
# cio_ignore -r 0.0.f500,0.0.f501,0.0.f502
```

3. Use the **znetconf** utility to sense and list candidate configurations for network devices:

```
# znetconf -u
Scanning for network devices...
Device IDs                Type      Card Type      CHPID Drv.
-----
0.0.f500,0.0.f501,0.0.f502 1731/01 OSA (QDIO)      00 qeth
0.0.f503,0.0.f504,0.0.f505 1731/01 OSA (QDIO)      01 qeth
0.0.0400,0.0.0401,0.0.0402 1731/05 HiperSockets 02 qeth
```

4. Select the configuration you want to work with and use **znetconf** to apply the configuration and to bring the configured group device online as network device.

```
# znetconf -a f500
Scanning for network devices...
Successfully configured device 0.0.f500 (enccw0.0.f500)
```

5. Optionally, you can also pass arguments that are configured on the group device before it is set online:

```
# znetconf -a f500 -o portname=myname
Scanning for network devices...
Successfully configured device 0.0.f500 (enccw0.0.f500)
```

Now you can continue to configure the **enccw0.0.f500** network interface.

Alternatively, you can use **sysfs** attributes to set the device online as follows:

1. Create a **qeth** group device:

```
# echo read_device_bus_id,write_device_bus_id,data_device_bus_id >
/sys/bus/ccwgroup/drivers/qeth/group
```

For example:

```
# echo 0.0.f500,0.0.f501,0.0.f502 >
/sys/bus/ccwgroup/drivers/qeth/group
```

2. Next, verify that the **qeth** group device was created properly by looking for the read channel:

```
# ls /sys/bus/ccwgroup/drivers/qeth/0.0.f500
```

You can optionally set additional parameters and features, depending on the way you are setting up your system and the features you require, such as:

- **portno**
- **layer2**
- **portname**

3. Bring the device online by writing **1** to the online **sysfs** attribute:

```
# echo 1 > /sys/bus/ccwgroup/drivers/qeth/0.0.f500/online
```

4. Then verify the state of the device:

```
# cat /sys/bus/ccwgroup/drivers/qeth/0.0.f500/online
1
```

A return value of **1** indicates that the device is online, while a return value **0** indicates that the device is offline.

5. Find the interface name that was assigned to the device:

```
# cat /sys/bus/ccwgroup/drivers/qeth/0.0.f500/if_name
enccw0.0.f500
```

Now you can continue to configure the **enccw0.0.f500** network interface.

The following command from the **s390utils** package shows the most important settings of your **qeth** device:

```
# lsqeth enccw0.0.f500
Device name                : enccw0.0.f500
-----
card_type                  : OSD_1000
cdev0                      : 0.0.f500
cdev1                      : 0.0.f501
cdev2                      : 0.0.f502
chpid                      : 76
online                     : 1
portname                   : OSAPORT
portno                     : 0
state                      : UP (LAN ONLINE)
priority_queueing          : always queue 0
buffer_count               : 16
layer2                     : 1
isolation                   : none
```

10.6. PERSISTENTLY ADDING A QETH DEVICE

To make your new **qeth** device persistent, you need to create the configuration file for your new interface. The network interface configuration files are placed in the **/etc/sysconfig/network-scripts/** directory.

The network configuration files use the naming convention **ifcfg-*device***, where *device* is the value found in the **if_name** file in the **qeth** group device that was created earlier, for example **enccw0.0.09a0**. The **cio_ignore** commands are handled transparently for persistent device configurations and you do not need to free devices from the ignore list manually.

If a configuration file for another device of the same type already exists, the simplest way to add the config file is to copy it to the new name and then edit it:

```
# cd /etc/sysconfig/network-scripts
# cp ifcfg-enccw0.0.09a0 ifcfg-enccw0.0.0600
```

To learn IDs of your network devices, use the **lsqeth** utility:

```
# lsqeth -p
devices                CHPID interface        cardtype        port
chksum prio-q'ing rtr4 rtr6 lay'2 cnt
-----
0.0.09a0/0.0.09a1/0.0.09a2 x00  enc0.0.09a0  Virt.NIC QDIO  0  sw
always_q_2 n/a  n/a  1  64
0.0.0600/0.0.0601/0.0.0602 x00  enc0.0.0600  Virt.NIC QDIO  0  sw
always_q_2 n/a  n/a  1  64
```

If you do not have a similar device defined, you must create a new file. Use this example of **/etc/sysconfig/network-scripts/ifcfg-0.0.09a0** as a template:

```
# IBM QETH
DEVICE=enccw0.0.09a0
BOOTPROTO=static
IPADDR=10.12.20.136
NETMASK=255.255.255.0
ONBOOT=yes
NETTYPE=qeth
SUBCHANNELS=0.0.09a0,0.0.09a1,0.0.09a2
PORTNAME=OSAPORT
OPTIONS='layer2=1 portno=0'
MACADDR=02:00:00:23:65:1a
TYPE=Ethernet
```

Edit the new **ifcfg-0.0.0600** file as follows:

1. Modify the **DEVICE** statement to reflect the contents of the **if_name** file from your **ccw** group.
2. Modify the **IPADDR** statement to reflect the IP address of your new interface.
3. Modify the **NETMASK** statement as needed.
4. If the new interface is to be activated at boot time, then make sure **ONBOOT** is set to **yes**.
5. Make sure the **SUBCHANNELS** statement matches the hardware addresses for your qeth device.
6. Modify the **PORTNAME** statement or leave it out if it is not necessary in your environment.
7. You can add any valid **sysfs** attribute and its value to the **OPTIONS** parameter. The Red Hat Enterprise Linux installation program currently uses this to configure the layer mode (**layer2**) and the relative port number (**portno**) of **qeth** devices.
The **qeth** device driver default for OSA devices is now layer 2 mode. To continue using old **ifcfg** definitions that rely on the previous default of layer 3 mode, add **layer2=0** to the **OPTIONS** parameter.

/etc/sysconfig/network-scripts/ifcfg-0.0.0600

```
# IBM QETH
DEVICE=enccw0.0.0600
BOOTPROTO=static
IPADDR=192.168.70.87
```

```

NETMASK=255.255.255.0
ONBOOT=yes
NETTYPE=qeth
SUBCHANNELS=0.0.0600,0.0.0601,0.0.0602
PORTNAME=OSAPORT
OPTIONS='layer2=1 portno=0'
MACADDR=02:00:00:b3:84:ef
TYPE=Ethernet

```

Changes to an **ifcfg** file only become effective after rebooting the system or after the dynamic addition of new network device channels by changing the system's I/O configuration (for example, attaching under z/VM). Alternatively, you can trigger the activation of a **ifcfg** file for network channels which were previously not active yet, by executing the following commands:

1. Use the **cio_ignore** utility to remove the network channels from the list of ignored devices and make them visible to Linux:

```
# cio_ignore -r
read_device_bus_id,write_device_bus_id,data_device_bus_id
```

Replace *read_device_bus_id*, *write_device_bus_id*, *data_device_bus_id* with the three device bus IDs representing a network device. For example, if the *read_device_bus_id* is **0.0.0600**, the *write_device_bus_id* is **0.0.0601**, and the *data_device_bus_id* is **0.0.0602**:

```
# cio_ignore -r 0.0.0600,0.0.0601,0.0.0602
```

2. To trigger the uevent that activates the change, issue:

```
# echo add > /sys/bus/ccw/devices/read-channel/uevent
```

For example:

```
# echo add > /sys/bus/ccw/devices/0.0.0600/uevent
```

3. Check the status of the network device:

```
# lsqeth
```

4. Now start the new interface:

```
# ifup encw0.0.0600
```

5. Check the status of the interface:

```
# ip addr show encw0.0.0600
3: encw0.0.0600: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc
pfifo_fast state UP group default qlen 1000
link/ether 3c:97:0e:51:38:17 brd ff:ff:ff:ff:ff:ff
inet 10.85.1.245/24 brd 10.34.3.255 scope global dynamic
encw0.0.0600
valid_lft 81487sec preferred_lft 81487sec
inet6 1574:12:5:1185:3e97:eff:fe51:3817/64 scope global
noprefixroute dynamic

```



```
valid_lft 2591994sec preferred_lft 604794sec
inet6 fe45::a455:eff:d078:3847/64 scope link
valid_lft forever preferred_lft forever
```

6. Check the routing for the new interface:

```
# ip route
default via 10.85.1.245 dev encw0.0.0600 proto static metric 1024
12.34.4.95/24 dev enp0s25 proto kernel scope link src 12.34.4.201
12.38.4.128 via 12.38.19.254 dev enp0s25 proto dhcp metric 1
192.168.122.0/24 dev virbr0 proto kernel scope link src
192.168.122.1
```

7. Verify your changes by using the **ping** utility to ping the gateway or another host on the subnet of the new device:

```
# ping -c 1 192.168.70.8
PING 192.168.70.8 (192.168.70.8) 56(84) bytes of data.
64 bytes from 192.168.70.8: icmp_seq=0 ttl=63 time=8.07 ms
```

8. If the default route information has changed, you must also update **/etc/sysconfig/network** accordingly.

10.7. CONFIGURING AN IBM Z NETWORK DEVICE FOR NETWORK ROOT FILE SYSTEM

To add a network device that is required to access the root file system, you only have to change the boot options. The boot options can be in a parameter file, however, since the **/etc/zipl.conf** no longer contains specification of boot records. The file that needs to be modified can be located using the following commands:

```
# machine_id=$(cat /etc/machine-id)
# kernel_version=$(uname -r)
# ls /boot/loader/entries/$machine_id-$kernel_version.conf
```

There is no need to recreate the initramfs.

Dracut, the **mkinitrd** successor that provides the functionality in the initramfs that in turn replaces **initrd**, provides a boot parameter to activate network devices on IBM z early in the boot process: **rd.znet=**.

As input, this parameter takes a comma-separated list of the **NETTYPE** (qeth, lcs, etc), two (lcs, etc) or three (qeth) device bus IDs, and optional additional parameters consisting of key-value pairs corresponding to network device sysfs attributes. This parameter configures and activates the IBM z network hardware. The configuration of IP addresses and other network specifics works the same as for other platforms. See the **dracut** documentation for more details.

The **cio_ignore** commands for the network channels are handled transparently on boot.

Example boot options for a root file system accessed over the network through NFS:

```
root=10.16.105.196:/nfs/nfs_root cio_ignore=all,!condev
rd.znet=qeth,0.0.0a00,0.0.0a01,0.0.0a02,layer2=1,portno=0,portname=OSAPORT
```

```
ip=10.16.105.197:10.16.105.196:10.16.111.254:255.255.248.0:nfs-server.subd  
omain.domain:enccw0.0.09a0:none rd_NO_LUKS rd_NO_LVM rd_NO_MD rd_NO_DM  
LANG=en_US.UTF-8 SYSFONT=latarcyrheb-sun16 KEYTABLE=us
```